Measuring Risk from Crop Portfolio Choices in Development Interventions *

Vincent Armentano†

October 30, 2026

Job Market Paper

Click here for the latest version.

Abstract

Development policy has promoted crop diversification for decades as a way to reduce smallholder farmers' exposure to risk. In spite of this, the standard measures of diversification, crop counts and concentration indices, cannot sign changes in portfolio risk. I make the long-standing analogy between crop choice and financial portfolio choice exact, estimating expected returns and the full covariance matrix of crop returns from plot-crop-level panel data collected for a cash transfer experiment in Southern Mali. Theory leaves the risk response to that transfer ambiguous: it could buy higher returns or lower risk. Farmers randomly assigned the grant reduced portfolio risk but did not sacrifice expected returns. Bounding risk aversion from where farmers sit relative to the efficient frontier, I value this gain at 5.2 to 12 percent of the transfer value per season. Standard count-based and specialization measures show significant effects but miss this improvement in efficiency entirely. The richest households in the sample improve risk-adjusted returns by expanding their portfolio breadth, consistent with a fixed cost of entering a new crop. In contrast, the poorest recipients do not expand their portfolios but reduce risk through reallocating land across crop types. This evidence points to a liquidity constraint on diversification rather than a shift in preferences for risk.

JEL codes: D24, G11, O12, O13, Q12, Q14

Keywords: cash transfers, modern portfolio theory, crop diversification, agricultural finance

*I am grateful to my dissertation co-chairs, Craig McIntosh and Tom Vogl, and committee members Sara Lowes and Paul Niehaus for their support and guidance. I would like to thank Sam Bazzi, Prashant Bharadwaj, Julie Cullen and Karthik Muralidharan as well as seminar participants at the Center for Effective Global Action and UC San Diego for their feedback. I am indebted to Chris Udry, Dean Karlan, Lori Beaman and Bram Thuysbaert for their generosity in sharing these data with me.

†University of California, San Diego. varmenta[at]ucsd.edu. Current draft and materials: vincentarmentano.com.

1 Introduction

How smallholder farmers cope with risk is a foundational question in development economics: uninsured risk pushes farmers to sacrifice expected income for safety (Rosenzweig and Binswanger, 1993), distorts investment away from profitable technologies (Morduch, 1995; Dercon, 1996), and transmits production shocks into consumption (Townsend, 1994; Kazianga and Udry, 2006). To manage their uninsured production risk, smallholders diversify their crop portfolios. Economists have been making the analogy between crop selection and the canonical portfolio selection problem in finance for over seventy years (Heady, 1952; Just and Pope, 1978), and yet the empirical literature still relies on simple proxies, such as the count of unique crop types, to measure portfolio risk. In this paper I provide a framework to make Modern Portfolio Theory (MPT) empirically tractable in applied agricultural settings by estimating expected returns and the full variance-covariance matrix of crop returns. This yields a measure of risk exposure for each household that accounts for all the strategic combinations of crops they cultivate. It also provides a benchmark of portfolio efficiency against which any intervention that changes what farmers grow (new seeds, credit, insurance, extension) can be evaluated. I then apply the framework to a large unconditional cash transfer experiment in Southern Mali to trace how farmers' portfolios respond (Beaman et al., 2023).

In this paper I estimate the covariance structure of crops at the level where planting decisions are actually made. With it, I can sign the risk change from any intervention that moves a household's crop mix, which is exactly where counts and concentration indices such as the Herfindahl-Hirschman Index (HHI) fall short. Existing efforts to understand diversification in agriculture have relied on these proxy measures (e.g., Di Falco and Chavas, 2009; Michler and Josephson, 2017; Kurosaki and Fafchamps, 2002; Makate et al., 2022). Among 74 economics papers specifically focused on measuring crop portfolio diversification since Heady (1952), 43% use the number of unique crops and 58% specialization indices; none computes the portfolio variance of households' observed crop portfolios (Figure A.1).[1] Prior to this paper, the covariance structure of crop returns has been applied in economics only to aggregates, at the district level (Allen and Atkin, 2022) and the county level (Levin and Vidart, 2026).

I exploit data from an unconditional cash transfer experiment in Southern Mali covering roughly 3,000 smallholder households across 110 villages in the Sikasso Region, observed over four survey waves from 2010 to 2017 (Beaman et al., 2023). An unconditional cash transfer program provides an unusually clear window into the way poverty constrains farmers' risk management choices. Most other agricultural experiments promote the take-up of a specific crop or technology, while a grant is general and does not alter the menu of crops available. In this setting agriculture is exclusively rain-fed, and with almost no access to formal finance

[1] The only farm-level implementations to date are two horticulture applications outside the economics literature (Paut et al., 2019; Burbano-Figueroa et al., 2022).

or insurance, the crop portfolio itself is the household's principal instrument for smoothing weather and price shocks. The experiment provided $140 ($153 in 2017 USD) to randomly selected treated households as a grant, about 12.3% of the control group's baseline expenditure on agricultural inputs. In theory, the response is ambiguous: the grant could raise risk-bearing capacity and push farmers toward higher-return crops. Alternatively, it could clear fixed costs of entering new crops and buy lower risk instead. First, I estimate the causal effect of the grant on portfolio risk and return, then the efficiency of the resulting portfolios, and finally I examine heterogeneity across the wealth distribution that reveals multiple mechanisms behind the average response.

My approach separates what the standard proxies bundle: how much risk a household bears and what return it earns for bearing it. The grant left expected returns unchanged but significantly reduced risk exposure. Portfolio standard deviation fell by $4.01 per hectare (1.3% of the $316 per hectare control mean). This combination led to a significant improvement in risk-adjusted returns as measured by the Sharpe ratio[2] (0.70% of the 95.0 control mean). Treated households spent roughly $31 more on production out of the $153 grant; scaling the risk reduction by that increase gives 0.23% less portfolio risk per 1% more spending.[3] Separating expected return from risk is what makes this improvement visible and the proxies economists have used record neither.

Portfolio theory offers a second lens on this improvement: efficiency. The efficient frontier is the set of portfolios that deliver the highest possible return at each level of risk; a household whose portfolio sits inside the frontier could earn more without bearing more risk (Markowitz, 1952). The grant moved farmers significantly closer to the frontier, a 3.0% reduction of the $49 per hectare control mean distance. As households' realized portfolio positions are constrained away from the frontier for reasons such as liquidity constraints, the closest spot on the efficient frontier gives us a way to estimate their preferences for risk absent constraints, an "ideal" point. This "ideal" point on the efficient frontier closest to each household did not move in response to the grant. Cash transfers enabled households to move toward the frontier without moving along it, a pattern that looks like the relief of a constraint rather than a change in preferences for risk.

The fact that relaxing credit constraints leads to more efficient portfolios without changing preferences for risk illustrates one important way crop portfolios differ from financial ones. In textbook portfolio theory, one's level of wealth plays no role in diversification. Efficient portfolio weights are the same at any scale, so a richer investor just holds more of the same mix. If relaxing credit constraints buys better diversification,

[2]As there is no risk-free asset available in my setting, the Sharpe ratio is defined simply as a portfolio's expected return divided by its standard deviation.

[3]This is a ratio of two treatment effects rather than a technological elasticity: it describes what this grant delivered through the spending it induced.

that suggests a cost of entering new crops was holding these portfolios at inefficient positions inside the frontier. Entry costs are easy to see in agriculture: each additional crop brings its own labor calendar, needs for chemical inputs and approach to land preparation. They also show up in spending: within a household, holding land and year fixed, a 1% increase in the number of crops cultivated is associated with 0.50% higher input expenditure. The finance literature met the same friction in transaction and membership fees and extended the theory to portfolios with a limit on the number of assets held (Chang et al., 2000); under such a limit, optimal portfolios can sit strictly inside the unconstrained frontier. The wealth gradient in the data matches this logic. Households in the richest wealth tercile are 41% more likely than the poorest to cultivate at least six unique crop types, and 94% more likely to cultivate at least eight. Constrained frontiers that limit the number of unique assets that can be held at once make the cost of these limits visible. Each cap on the number of crops shrinks the set of available portfolios toward the monocropping positions, reducing the amount of risk a household can diversify away.

Fixed costs in agriculture may cause responses to the grant to vary across levels of wealth. Heterogeneity by baseline wealth terciles shows both tails of the distribution reducing risk, but through different channels. Portfolio standard deviation falls by \$6.0 per hectare in the poorest tercile and by \$5.1 per hectare in the richest. The poorest households reduce risk without expanding their portfolios: land is reallocated as one crop substitutes for another, but the number of crops they hold does not rise. They likely cannot overcome the fixed costs of expanding their portfolios, but still their distances to both the unconstrained frontier and the relevant crop-count-constrained frontier fall significantly, which suggests a reallocation of land towards inherently safer crops. The richest households instead work through entry into new crops. They add on average 0.4 unique crop types, their Sharpe ratios improve, and the crops they bring in behave as hedges for the rest of the portfolio. Decomposing portfolio variance into a crop-selection channel and a hedging channel is consistent with this split: the poorest tercile's risk reduction shows up on the crop-selection margin, while the richest tercile's added crops behave as hedges that exploit the covariance structure.[4] A fixed cost of entering a new crop explains both responses at once. Cash relaxes a liquidity constraint, but the entry cost still binds for the poorest households, so they shuffle which crops they grow without expanding portfolio size. For the richest, it no longer binds, and they expand.

A lack of appropriate tools has largely prevented development economists from engaging with portfolio risk in recent research. Of the 327 agricultural risk experiments registered in the AEA RCT registry, only 5.8% pre-specify any outcomes for portfolio changes on the intensive or extensive margin.[5] I hope that with

[4] As the experiment was stratified on village rather than wealth, the differences between terciles are estimated too imprecisely to be tested directly.

[5] The AEA RCT Registry (socialscienceregistry.org), accessed July 2026: 12,514 registered trials, of which 327 concern risk in agriculture (title, abstract, keyword, and JEL screen); outcome presence was coded from only pre-specified primary and secondary outcomes.

this method for applying MPT to smallholder agriculture in hand, more researchers will explore how their interventions shift household responses holistically. To this point, I want to emphasize that nothing about the framework is specific to cash transfers or to Mali.[6] The required ingredients are estimated entirely from panel variation among untreated farmers, so the framework requires no experiment of its own and can be laid over any intervention, program placement, or policy change that alters what farmers grow.

What MPT does ask for is production revenues and costs measured at the crop-plot level within a household's farming season. Mechanically, the number of unique crops and Herfindahl-Hirschman Index are less data intensive to employ than applying MPT as they only require a crop list or land allocations from respondents while MPT requires knowing the crop properties alongside the land allocations. However, for researchers concerned with risk, the gains from measuring it properly will outweigh the cost of collecting the additional data. The data requirements can be relaxed if researchers are running experiments or doing analysis in a country where large panel surveys are already present, for example the Living Standards Measurement Survey-Integrated Surveys on Agriculture (LSMS-ISA) countries and in the near future the 50 countries that are participating in the World Bank's 50x2030 initiative to expand agricultural data collection. Researchers could calculate crop properties like expected return and the variance-covariance matrix of returns in a sub-sample of the LSMS-ISA data that represents their experimental sample well, and then just use the weights collected from the experimental sample to calculate portfolio risk and expected return values.

The results speak immediately to how diversification is measured. An agricultural-economics tradition extended mean-variance ideas to farm planning with calibrated programming models (Freund, 1956; Hazell, 1971; Collins and Barry, 1986; Blank, 1991). More recent panel-data work documents portfolio responses to drought with count-based measures. Asset-rich households in Malawi and Tanzania add crops (Makate et al., 2023) while resource-poor Ethiopian households shed them (Makate et al., 2022). Once a covariance matrix is in hand, counting crops is neither necessary nor sufficient for signing changes in risk. The crops treated farmers added happened to reduce portfolio risk, but the same count increase could as easily have raised it. This paper contributes modern panel data, causal identification, and an estimated covariance matrix to that agenda.

A foundational literature derives the risk-management mechanism from elicited risk preferences and relies on point-identifying the coefficient of relative risk aversion (γ) to interpret observed portfolio choice (Rosenzweig and Binswanger, 1993; Dercon, 1996). A structural strand instead runs the inversion in the other direction, treating observed crop portfolios as the revealed-preference data from which risk attitudes

[6] I employ the agricultural setting because the benefit is clear, but the benefits apply anywhere in the literature where specialization is being used as a proxy for risk. For example an application to roadside fruit sellers may show similar levels of specialization across the types of fruit they sell, while MPT could reveal that one's selection of fruit is truly much riskier than the others on a basis of profit per square meter of shelf space.

are recovered (Bar-Shira et al., 1997; Kurosaki and Fafchamps, 2002; Koundouri et al., 2009; Allen and Atkin, 2022), most recently for historical U.S. populations by Levin and Vidart (2026). Rather than point-identify a risk aversion parameter for each household, I follow Calvet et al. (2007)'s frontier-projection approach. I project each household onto its frontier in two directions: the position it would hold if it cared only about raising expected return, and the position it would hold if it cared only about lowering variance. Those two projections return a range of admissible γ values for each household rather than a single imputed one. Applying those bounds to the estimated treatment effects on portfolio mean and variance places the risk reduction at 5.2–12% of grant value per season (26–61% per marginal dollar the grant moved into agricultural production).

The Modern Portfolio Theory framework also helps us think about the impacts of other types of interventions that structurally shift risk-return trade-offs. Through the lens of my framework for example, the rainfall and weather index insurance interventions of Karlan et al. (2014) in Ghana, Mobarak and Rosenzweig (2013) in India, and Cai (2016) in China can all be read as introducing a new asset, an insurance asset, with strictly negative correlations to the insured crops and its own expected return. How insurance changes planting behavior can then be understood through the risk-return lens. Insurance relaxes the risk constraint and pushes farmers toward riskier, higher-return production. For example Karlan et al. (2014) shows farmers narrow their portfolios to specialize in high-intensity maize. Cash pushes in the opposite direction in Mali because the binding constraint is different: risk-bearing capacity there, the cost of entering new crops here.

Finally, my work adds to a wide and varied literature exploring how cash transfers impact the household-firm. Previous work has documented robust effects on consumption, savings, and asset accumulation (Haushofer and Shapiro, 2016; Egger et al., 2022; Bandiera et al., 2017). The contribution here is not that cash unlocks investment, which Beaman et al. (2023) established with this same experiment, but a new dimension of how to understand what household-firms did with the loosened constraint. I show they exploit crop covariances to reduce portfolio variance at constant expected return, the revealed-preference signature of a previously binding diversification constraint. In this setting, the cash was spent on safety and not risk-taking.

2 Methodology

2.1 Application of MPT & Limitations of Financial Asset Analogy

MPT in its modern form mostly exists in different variations of the Capital Asset Pricing Model (CAPM), which is a general equilibrium model originating in Sharpe (1964). For this analysis I rely only on the

model's partial equilibrium elements, which are 1) the expected return of each asset, 2) the variance of each asset's returns and 3) the covariance matrix of asset returns. These are all that is required to model the diversification benefit of a farmer's crop portfolio choices. As everything is done in profit per hectare, factors that move both costs and revenues across crop types will be accounted for by the estimated covariance terms. This normalization also means that farms of different sizes can be compared more easily. Without the estimation of covariance terms, I would miss the impacts of any strategic crop combinations cultivators are making, which are critical to understanding their exposure to risk.

2.1.1 Removal of the Risk-Free Rate and Capital Market Line

Two further elements of CAPM, the risk-free rate and the capital market line, do not carry over to this setting; Figure A.2 illustrates both. The risk-free rate is the return on an asset like a bond, which bears no risk and a commensurately low return relative to available risky assets. Note that the blue curve, referred to as the efficient frontier (EF), marks the boundary of what portfolio positions are feasible given investments in only risky assets. The line extending out from the risk-free asset's location to be tangent to the EF is called the capital market line (CML) and represents portfolios consisting of mixes between the risky assets and the risk-free asset. As return is an economic good and risk is an economic bad, the CML weakly dominates all points along the efficient frontier for individuals looking to maximize return and minimize variance, regardless of their taste for risk. Individuals then organize themselves along the CML according to their preferences for risk. The most risk-averse individuals combine the risk-free asset with the tangent portfolio and sit on the CML below it, while those willing to borrow at the risk-free rate leverage into the tangent portfolio and sit above it.

However, this important feature of CAPM does not exist without the risk-free rate, which is an asset class that has no equivalent in smallholder agriculture. While land rental rates, or wages that could be earned from laboring if farmers abstained from participating in agriculture could hypothetically stand in for the risk-free asset, the land markets and labor markets in smallholder settings are often incomplete and fail to be reliable sources of returns the way the risk-free asset needs to be. Without the risk-free asset and capital market line, risk-averse investors who would otherwise lend at the risk-free rate must instead choose portfolios of risky assets with relatively lower returns along the EF. Individuals more tolerant of risk will inhabit the portion of the curve with higher levels of return and higher levels of risk. Figure A.3 illustrates this point with indifference curves representing preferences coming from the mean-variance utility function, $U(\mu, \sigma) = \mu - \frac{\phi}{2}\sigma^2$ where ϕ is a risk aversion parameter governing the trade-off between risk and return.

The EF represents both the result of maximizing return given a set level of risk and minimizing risk for a desired level of return. Portfolios along this frontier are deemed "efficient" because, at their level of

expected return, they have removed all the idiosyncratic risk that can be diversified away through portfolio choice. This measure of efficiency will be very important to my analysis of risk taking behavior as it provides a grounding and benchmark for discussing how constrained individuals are. More detail on this is provided in Section 3. While this analysis relies only on the foundational elements of MPT, treating crops as one would financial assets, there are important limitations in this analogy involving returns to scale, interaction effects from intercropping, technological complementarities across crops, the impossibility of short sales, and the indivisibility of crop allocations. I address each of these limitations directly. Farmers in my sample do not exhibit economies of scale (Figure I.1) and I have too few repeat instances of cultivation to account for intercropping complementarities (Appendix I.2). The indivisibility of crop allocations is taken up directly in Section 5.6, where I explore a cardinality-constrained version of the MPT problem and illustrate how a per-crop fixed-cost threshold reshapes the concept of the efficient frontier. The cardinality-constrained discussion in Section 5.6 is descriptive rather than a fully structural utility-maximization solution. The interpretation of the associated results therefore relies on the qualitative consistency of the supporting evidence rather than on a quantitative match between a calibrated model and the empirical treatment effects. A full discussion of these dimensions and their implications for applying MPT in other settings is provided in Appendix I.

2.1.2 No Short Sales and the Bounded Long-Only Frontier

A third element of the financial analogy that does not carry over is one of short sales. An investor may hold a negative position in an asset, funding a larger position in another; but a farmer cannot plant negative hectares of sorghum. Portfolio weights here are therefore non-negative by construction, $w_i \geq 0$ with $\sum_i w_i = 1$, and this is a physical property of land allocation rather than a modeling convenience.

The restriction has a consequence that matters for how I measure efficiency. With non-negative weights the feasible set is the convex hull of portfolios containing any number of crops, so the highest attainable expected return is that of the single highest-returning crop, and the efficient frontier terminates there. To obtain the maximum return on the efficient frontier, the household must hold one crop and nothing else. This is mechanically also a function of how weighted averages work. Permitting short sales would remove that endpoint, since a farmer could exceed any single crop's return by taking a negative position in a lower-return crop to fund a larger position in a high-return one, and the upper branch would extend without bound. The distance from a household's position to the positive weights only frontier, which is the efficiency measure used throughout this paper, is the relevant measure given my setting.

The two ends of the frontier are asymmetric in a way that is worth making explicit, because the asymmetry is my paper's thesis in miniature. The upper end is pinned to a single crop by the argument just given, but the lower end is not. The global minimum-variance portfolio combines several imperfectly correlated

crops and attains a standard deviation strictly below that of the least volatile crop available, so the frontier extends past every individual crop marker at the low-risk end. That gap between the minimum-variance portfolio and the safest single crop is exactly the risk reduction that a crop count cannot correctly capture the sign of, since both options are described by a small integer, and only the portfolio calculation values the difference in risk between them.

2.1.3 Handling Intensification in the Analogy

There is no other product that an investor could buy to increase the expected return of an asset in their portfolio. In agriculture, however, farmers can change the intensity with which they cultivate a crop by applying more fertilizer, hiring additional labor, or investing in land preparation. These activities can alter both the expected return and the variance of that crop's profit per hectare. This means that within a single crop label, there can exist effectively multiple distinct production technologies, each yielding different return distributions. A farmer growing maize with purchased fertilizer and hired labor faces a different risk-return profile than one growing maize under traditional low-input methods, even though both are nominally cultivating the same crop. If these distinct technologies are pooled into a single asset class, the estimated return moments will average over heterogeneous regimes and mischaracterize the risk properties of both.

To address this difference, I apply Hartigan's Dip Test of unimodality (Hartigan and Hartigan, 1985) to the distribution of input costs per hectare for each crop type. I split crops whose cost distributions reject unimodality (Figure A.4) into high-intensity and low-intensity variants at the control-group mean cost per hectare, each treated as a separate crop type with its own expected return and variance for the rest of the analysis. I want to clarify two points about what this procedure does and does not establish. The dip test's null hypothesis is unimodality, so a rejection indicates only that a single mode does not describe the cost distribution; it does not identify how many regimes are present, and cutting at a single threshold imposes two where there could be more. For maize, a unimodal distribution is rejected and the two new crop definitions are split along the control-group mean cost per hectare, \$210/HA for maize. Through this method, input intensity is now able to be respected so the covariance matrix can distinguish high- from low-input cultivation of the same crop rather than averaging over both. Rice, shown in contrast as a crop whose cost distribution is unimodal ($p = 0.86$), is treated as a single asset class. Appendix C.2 reports treatment effects on the split and merged versions of each affected crop, and Section 6 reports the headline effect under both conventions; the risk result does not depend on the choice. Appendix C.2 shows that share-level treatment effects on the split (high-intensity, low-intensity) and merged versions of each affected crop are mutually consistent, so the dip-test split is retained as the primary specification because it identifies real input heterogeneity rather than because the headline result depends on it.

2.2 Estimating Expected Return from Profit Per Hectare Measurements

Profits are notoriously hard to estimate in smallholder agriculture given the self-reported nature of the data, measurement error and the true heterogeneity in input costs, cropping arrangements, sale prices and output quality experienced by people working as smallholder farmers. In agriculture it's not uncommon for simple revenue to be used as a proxy for income. It's much more salient for respondents to provide researchers and is more easily attributable to specific crops than the expenditures incurred in production, which are almost always in part shared across crops, like fertilizer used on an intercropped plot. Only looking at revenue unfortunately hides the heterogeneity in costs that exists across crop types and is likely playing a role in crop selection. In the data, what could have seemed like a very good season for maize could have truly been mediocre after accounting for extensive fertilizer expenditures, while cassava, requiring little additional inputs and land preparation to cultivate, could have had underwhelming revenues in levels but then looked more profitable and attractive after accounting for its few costs. As my data contain both revenues at the crop-plot level and input expenditures at the plot level: for sole-cropped plots I can create crop specific profit per hectare values from the data without any additional assumptions. Because most input expenditures are recorded only at the plot level though, intercropped plots require a cost-allocation rule. I parse these costs across crops according to their revenue share on the plot: a crop that earned twice as much as its counterpart on a two-crop intercropped plot receives twice the costs of the other when creating profit per hectare.[7] I begin my analysis by creating a plot-level profit per hectare measure, representing revenue less all cash expenditures on hired labor, land preparation, chemicals and seeds used on the plot divided by the area cultivated in hectares.

The two main types of systematic risk at play in agricultural diversification are weather and price risk. Weather risk directly impacts yields, which on an aggregate level affects prices through the supply side of the market. The relationship between yields and prices brought on by weather is often referred to as a natural form of insurance for smallholder farmers (Newbery and Stiglitz, 1981). Good weather brings large harvests and low prices; bad weather brings small harvests and higher prices, partially offsetting the shortfall in revenue relative to expectations. The relationship of local weather and local prices is a function of how locally determined prices actually are though and the recent piece Allen and Atkin (2022) illustrates just this as they show that integration of rural markets to the world economy makes sale prices less responsive to local weather conditions. Appendix C.10 separates the two in the data: within farmer-crop cells, yield

[7] The same revenue-share convention also governs how plot area is allocated across constituent crops on an intercropped plot, so that the per-crop hectare weights w_c always sum to one within a household-year as the MPT framework requires. Appendix C.1 details the alternative cost and area allocation conventions considered, the empirical extent of disagreement between them, and the cross-allocation robustness specifications. The headline portfolio-SD treatment effect survives the most direct area-allocation robustness in Appendix D.3.

variation accounts for 0.98 of the variance of log revenue per hectare, price variation for 0.04, and their covariance, where natural insurance would appear, is small and negative at -0.02. Prices are also going to be a function of transport costs for traders who aggregate harvests across space and then sell wholesale to exporting warehouses. My data record each respondent's sale prices and buyer type but to measure harvests at a common price while still allowing for spatial variation, I use the median sale price to traders for each crop type in each community to value harvest revenues. [8] When no sales were made to traders in a community, I take the sample-wide median sale price to value harvests. This fallback is common, as most production in this setting is consumed rather than sold: 74% of plot-crop harvests involve no own sale, and in 41% of community-round-crop cells no household sold the crop at all. Those cells are small, however, so the fallback reaches far fewer harvests than the cell count suggests: 21% of plot-crop harvest instances (47,817) are valued at the sample-wide median, and the remainder at their own community's. Because the same valuation rule applies identically to treatment and control households, this affects levels of the return measure rather than treatment comparisons. Linearly, these two dimensions of systematic risk combine in profits, which makes it much more attractive for me to employ profits when examining systematic risk than revenues.

When applying this framework to evaluate the effects of an intervention, the crop moments must be estimated from observations whose production decisions are not confounded by the treatment itself; otherwise the estimated return distribution could reflect treatment-induced production behavior changes rather than the underlying properties of each crop: so I limit the sample used for moment estimation to the control group. And to ensure I am capturing a central tendency of profit per hectare for a given crop, I further limit the sample used for my estimation to farmers who cultivated a crop in at least 2 seasons. This lets me pull both within-farmer variation and across-farmer variation out of the distribution depending upon the moment of interest.

To create the average return of a crop, I first remove the within-farmer variation by finding each farmer's average return for the given crop across the panel. This step removes the temporal variation from sources like weather and price variation. From these farmer-level averages, I next take the average across farmers who cultivated the crop, which removes the across-farmer variation. The resulting average value is what I use as the crop's expected return per hectare. This value captures the profit per hectare for this crop for an average farmer, in the average year observed in my sample. Figure A.5 illustrates this construction for two

[8] An alternative approach uses median *consumption* prices (what households report paying when purchasing food) rather than farm-sale prices to value non-sold production (De Magalhães and Santaeulàlia-Llopis, 2018; Rodríguez-Sala, 2026). Consumption prices may better capture the shadow value of subsistence crops, as sale prices are typically measured post-harvest when supply is abundant and prices are low. In this sample, the distinction is unlikely to affect treatment effects (same measurement in treatment and control), but it may modestly understate the true return for crops that are primarily consumed rather than sold. Valuing all production at market prices, whether sold or self-consumed addresses the alternative explanation that observed multicropping reflects a household preference for dietary variety rather than a risk management strategy. Any increase in utility from consuming a more varied basket is monetized through the market-price willingness-to-pay equivalent, so the documented portfolio-SD reduction is identified in income-variance space rather than dietary-breadth space.

contrasting crops: each dot is one farmer's average profit per hectare across the years they cultivated the crop, and the vertical bar marks the across-farmer mean that becomes $\bar{R}_c$ in the analysis. To calculate the Expected Return of a cultivator's portfolio, I then simply multiply the vector of crops' returns, $\vec{R}$, against a vector of hectare-weights detailing the shares of land a cultivator allocates to each crop.

2.3 Estimating the Variance of Each Crop Type and Covariances across crops

To estimate the variance of a crop's returns, I again start by limiting my sample to all instances where the same cultivator was observed growing the same crop in at least two farming seasons. This ensures that the within-farmer time variation plays a central role in the analysis. I then demean returns within each farmer-crop combination. I am then left with residuals centered around zero, showing deviations from that cultivators' own average profit per hectare values for the crop. A cultivator who grows the crop on more than one plot in a season therefore contributes more than one deviation, which is the case for 26% of farmer-crop-years, and it is why the average number of deviations per farmer-crop cell can exceed the 4 survey waves. These residuals are the within farmer variation that shows how the crop performed on the same farm across different states of nature. If these residuals are all small in magnitude, it highlights that the farmer will have a fairly good idea what they'll earn when they cultivate this crop across different states of nature and vice versa. This is what distinguishes low risk from high risk crops in my analysis. I then square these deviations across all farmers and take their sum, dividing by the residual degrees of freedom, $N_c - G_c$, where N_c is the number of deviations and G_c the number of farmer-crop cells contributing them (Appendix C.3). These variance estimates will then make up the diagonal entries of the variance-covariance matrix I use to evaluate portfolio variance. Figure A.6 shows the resulting density of these residuals for the same two crops as Figure A.5: cotton's much wider distribution ($\sigma = 729$/HA) relative to sorghum's ($\sigma = 294$/HA) is the cross-crop heterogeneity in idiosyncratic risk that the diagonal of $\mathbf{C}$ records. [9]

With this method of using residuals to measure variance, I face the threat of measurement error being present in my estimates of variance for each crop. Measurement error of the same scale across crop types inflates every crop's variance in levels but leaves their rank order unchanged. What would threaten their rank order is error that differs systematically across crops, large enough for one crop's inflation to overcome another crop's genuinely larger variance. Appendix C.3 tests this directly, following the within-farmer decomposition strategy of Gollin and Udry (2021). That decomposition recovers an upper bound for measurement error rather than a point estimate, since within-farmer variation also contains genuine late-season shocks and land-

[9] Note that the dispersion in Figure A.6 is being driven by within-farmer variation, while the dispersion seen in Figure A.5 is purely across-farmer variation, as the within-farmer variation has been removed from all the dots making up the observations in Figure A.5. This does well to illustrate that the across-farmer variation is not employed in the construction of the variance-covariance matrix.

quality differences these data cannot separate from mismeasurement. So bounded, the error share runs from 13% to 46% across the 6 major crops (mean 22%, coefficient of variation 0.54): under half for every major crop, though not uniform. Uniformity is not what the ranking interpretation requires. What it requires is that the observed ordering of crop variances survive the correction that differential mismeasurement would impose, and 93% of major-crop pairs do (Appendix C.3); because treatment is randomized and the survey instrument is identical across arms, the error is also balanced between arms within every crop.

The covariance terms are estimated one level up, at the farmer-year. I average each farmer's deviations for a crop within a season before computing the pairwise correlations, and consider only the farmers who cultivated a given pair of crops across at least two farming seasons. The diagonal of $\mathbf{C}$ accordingly carries the dispersion across a farmer's plots within a season while the off-diagonal does not, which is also what makes the within-season, across-plot comparison in Appendix C.3 usable as a bound on the measurement-error share of that diagonal. With the farmer-year level observations of deviations used above to create the variance terms, I estimate covariance terms by looking at how the residuals covary across time. Figure A.7 illustrates what those farmer-year deviations look like for two crop pairs: groundnut (low) and sorghum comove at $\rho = 0.50$, while groundnut (low) and cotton comove at only $\rho = 0.29$. Neither pair is negatively correlated, which will be the norm in agriculture given the shared factors in the production functions of each crop. Diversification value in this setting will come chiefly from imperfect comovement rather than from true hedges. It is the farmer-level object that governs the risk an individual household bears, so $\mathbf{C}$ is estimated there rather than at the community-year level, which would amplify systematic comovement and wash out the idiosyncratic (Appendix G.3). The returned covariances take into account the strength of the relationship and the relative riskiness of each of the crops. The exact relationship between covariance and correlation coefficients is given in Equation 1, where $\phi_{A,B}$ is the correlation term between returns of crop A & B and σ_X represents the standard deviation of crop X. I report a matrix of correlation coefficients to show that there are a range of crops that don't covary strongly and this is where opportunities for diversification exist for smallholder farmers against the counterfactual that all crops are perfectly correlated in their returns. I collect these values in the variance-covariance matrix $\mathbf{C}$. Because each entry is estimated from the farmers who grow the crop or crop pair in question, $\mathbf{C}$ is assembled from partially overlapping samples and positive semi-definiteness is not guaranteed by construction. It holds here: the smallest eigenvalue is 2,366 and the condition number is 629 (36 for the correlation matrix), so no nearest-positive-definite projection is applied anywhere in the paper.

$$\sigma_{A,B} = \phi_{A,B} \cdot \sigma_A \cdot \sigma_B \tag{1}$$

2.4 Selection into Cultivation and the Scope of the Moment Estimates

A concern for both objects is that $\vec{R}$ and $\mathbf{C}$ are only estimated from observed cultivation choices: if households with different productivity levels systematically sort into different crops, the estimated crop-specific moments may reflect the productivity of the underlying cultivators rather than properties of the crops themselves. Observable baseline household characteristics (farm size, agricultural assets, household assets, plough ownership, fertilizer use, and demographic composition) do correlate with which crops a household cultivates (Appendix J, Figure J.1). But the relevant question is whether or not this sorting is severe enough to systematically bias the moment estimates, and three pieces of evidence suggest it is not. First, within every major crop's cultivator pool, farm sizes span the full baseline distribution. For rice and high-intensity groundnut (the high-variance crops most likely to be flagged as candidates for capability-based selection), cultivators are drawn 15–19% from the smallest-farm quintile and 18–24% from the largest, close to a uniform 20% benchmark (Appendix Figure J.3). Within each crop's data the moments therefore average across farmers from across the productivity distribution. Second, if capability-based sorting were severe, the menu of crops would partition into capability-based clusters with limited cross-cluster cultivation. A pairwise Jaccard cultivation-overlap matrix (Appendix Figure J.2) shows no such block structure: the median pair of crops shares 10% cultivator overlap, and there is no clean partition of the menu into capable and less-capable subsets. Third, the cross-sectional ordering of σ_c is highly stable when residuals are demeaned at the community $\times$ year level rather than the farmer level alone (Spearman rank correlation $\rho = 0.985$; Appendix J.3), confirming that geographic-shock heterogeneity is not driving the variance ranking. Cotton is the one exception among major crops, with only 5.8% of its cultivators in the smallest farm-size quintile, a pattern consistent with the labor and processing requirements of cotton production. I flag this asymmetry directly and report robustness specifications excluding cotton in Appendix G.1.

I emphasize that the paper's causal claims are identified within the randomized experimental sample rather than extrapolated to a representative Malian farmer. The vector $\vec{R}$ and the matrix $\mathbf{C}$ are intended to characterize the menu of crop opportunities that this sample's farmers actually face when choosing their portfolios; the welfare analysis that follows is computed against that menu. Only having four survey waves also bounds how much of the state space these estimates observe; Appendix D assesses the resulting stability of the covariance estimates using the different agro-ecological zones present in the seven-wave Uganda LSMS-ISA panel. It finds that frequently observed crop-pairs will have stable estimated correlation coefficients while less frequently observed pairs are less likely to be accurately estimated in a four wave panel.

2.5 Estimating the Portfolio Properties

After obtaining the estimates of returns across crops, $\vec{R}$ and the variance-covariance matrix of crop returns $\mathbf{C}$, creating the desired Expected Return and variance values of a cultivator's portfolio can be done with simple matrix multiplication. Multiplying each of these objects by vectors of weights that cultivators have assigned to each crop returns the two values of interest: see Equations 2 & 3.

$$\mathsf{E}[R] = \vec{w}\,' \times \vec{R} \tag{2}$$

$$\text{Var} = \vec{w}' \times \mathbf{C} \times \vec{w} \tag{3}$$

2.6 Decomposing Variance Changes into Self-Insurance and Diversification Channels

The portfolio variance in Equation 3 can be separated into two additive components that have distinct economic interpretations:

$$\text{Var} = \underbrace{\sum_{i} w_i^2 \sigma_i^2}_{\text{Diagonal}} + \underbrace{2 \sum_{i<j} w_i w_j \text{Cov}_{ij}}_{\text{Off-diagonal}} \tag{4}$$

The diagonal terms capture the weighted sum of individual crop variances. A household that allocates land away from a high-variance crop toward a lower-variance crop will reduce their portfolio variance holding how the crops' returns covary fixed. This channel represents *self-insurance*: the household accepts lower expected returns in exchange for lower volatility by selecting inherently safer crops. In contrast, the off-diagonal terms capture the contribution of return covariances between crop pairs. The off-diagonal terms lower portfolio variance whenever crop returns are less than perfectly correlated ($\rho_{ij} < 1$): two crops whose returns move with $\rho_{ij} = 1$ are no safer together than either alone, and any departure from lockstep is a gain. Negative correlation makes the gain larger but is not required.

This channel represents *diversification* and hedging in the MPT sense: strategic crop combinations that exploit imperfect correlations to hedge risk.

Equation 4 identifies the two channels but does not itself yield outcome variables: its terms are in squared units, scale with portfolio size, and are not comparable across households facing different crop menus. I therefore construct two ratio measures, each an exact measure of one of the two channels rather than a normalization of the corresponding term. The distinction matters for how the two measures should be read: they are constructed to speak to the economic channels that the additive terms in Equation 4 represent, not to be the above component terms re-scaled. Neither ratio should be expected to reproduce the sign or

magnitude of its corresponding summand. The first asks how much the household's crops, once selected, affect portfolio variance from combining imperfectly; the second asks whether the crops a household selects are intrinsically safer than the benchmark its community offers.

2.6.1 Diversification Ratio

The first measure compares the portfolio's actual standard deviation to a counterfactual standard deviation recalculated where all crop returns are perfectly correlated:

$$\text{DR} = \frac{\sigma_p|_{\rho_{ij}=1}}{\sigma_p} \times 100 \tag{5}$$

The numerator, $\sum_i w_i \sigma_i$, is the portfolio standard deviation that would be obtained if all pairwise correlations equaled one: a "no-diversification" benchmark (Choueifaty and Coignard, 2008). The gain this measure records requires only that correlations fall short of one, not that any of them be negative. Choueifaty and Coignard's own illustration uses three assets whose pairwise correlations are 0.9, 0.1 and 0.1, all strictly positive, and the most-diversified portfolio still departs sharply from equal weights. That distinction matters in this setting, where the estimated correlation matrix contains almost no negative entries: the hedging value available to these households comes from the spread in positive comovement rather than from true offsets. Under perfect correlation there is no diversification benefit to holding multiple crops; the portfolio's risk is simply the weighted sum of individual crop risks. The denominator is the actual household estimated portfolio standard deviation, which reflects any risk reduction achieved through imperfect correlations.

A diversification ratio of 100 is the floor of the measure, and Choueifaty and Coignard (2008) show it is attained exactly when the portfolio is equivalent to a single-asset holding. A monocropping household scores 100 with no correlation entering the computation; perfectly correlated returns across several crops reach the same floor by a different route. Values above 100 therefore indicate both that the household holds more than one crop and that those crops' returns are imperfectly correlated. Values above 100 also indicate that imperfect correlations are actively reducing the household's risk below the no-diversification benchmark, with higher values indicating greater risk reduction from diversification. Unlike the selection ratio I describe below, which will measure the riskiness of crop selection relative to a profit-maximizing benchmark, the diversification ratio measures how much total risk has been reduced relative to a world without diversification opportunities.

Choueifaty and Coignard (2008) show that when expected excess returns across assets are proportional to their volatilities, maximizing the diversification ratio is equivalent to maximizing the portfolio Sharpe ratio, and the most-diversified portfolio coincides with the tangency portfolio. That proportionality is close

to an empirical fact here: the cross-crop correlation between expected return and standard deviation is 0.95 (Spearman 0.93, Figure A.8). On this menu the diversification ratio is therefore approximately the objective a Sharpe-maximizing household would hold, which is why a treatment-induced increase in it carries welfare content rather than describing weight dispersion alone.

2.6.2 Selection Ratio

The second measure benchmarks the household's linear-weighted average crop standard deviation against the standard deviation of the highest-return crop available:

$$\text{Selection Ratio} = \frac{\sigma_{\max}}{\sum_i w_i \sigma_i} \times 100 \tag{6}$$

where $\sigma_{\max}$ is the standard deviation of the crop with the highest expected return in the household's community and w_i is the household's acre share on crop i. The same risk-return proportionality underpins this measure: because higher-return crops are also the riskier ones, benchmarking a household's weighted crop risk against the highest-return crop available is informative about self-insurance. A selection ratio above 100 indicates that the household's land-share-weighted average crop standard deviation is lower than the standard deviation of the most profitable crop available to them; they are accepting lower expected returns in exchange for cultivating land in inherently less volatile crops. The linear weighting in the denominator distinguishes this measure from the squared-weight diagonal $\sum_i w_i^2 \sigma_i^2$, which mechanically falls when weights spread across crops even when the underlying portfolio of crops is held fixed; the linear-weighted formulation isolates the genuine crop-selection channel.[10] The linear weighting also keeps this measure clear of a measurement-error term that the squared-weight diagonal carries: observed portfolio variance is inflated in proportion to portfolio concentration, so a treatment that spreads weight across crops mechanically lowers the squared-weight object even absent any change in true risk (Appendix C.4).

The denominator $\sum_i w_i \sigma_i$ is the same linear-weighted-SD object that anchors the Choueifaty-Coignard diversification ratio numerator, providing a direct comparability between the self-insurance and hedging measures. Higher values indicate a greater gap between the riskiness of the profit-maximizing crop and the riskiness of the household's actual crop mix, reflecting stronger self-insurance through crop selection. Because the numerator is a community-level benchmark, cross-household variation in the selection ratio is driven entirely by portfolio composition: which crops a household grows and how it allocates land across them.

[10] The squared-weight diagonal is itself a legitimate object of interest (it is the literal diagonal contribution to the total portfolio variance $\sigma_p^2 = \sum_i w_i^2 \sigma_i^2 + 2 \sum_{i<j} w_i w_j \text{Cov}_{ij}$) and treatment effects on it are reported in Appendix G.2 (Table G.3). The interpretation of changes in the squared-weight diagonal is, however, ambiguous: a treatment-induced reduction can reflect either selection toward lower-variance crops, mechanical weight spreading across more crops, or both, and these channels are not separable from the squared-weight aggregate alone.

This measure mirrors the logic of the Diversification Ratio introduced by Choueifaty and Coignard (2008): both compare a household's realized portfolio risk to a single-asset benchmark, but while the Diversification Ratio captures the hedging benefit of imperfect correlations, the selection ratio captures the self-insurance benefit of selecting lower-variance crops.

Read together, the two measures step through one chain. The selection ratio sets the risk of the most profitable crop available against the household's linear-weighted average crop risk; the diversification ratio sets that same average against the risk the household actually bears. Between them the two steps span the distance from $\sigma_{\max}$ down to σ_p: the first is the part of that distance closed by choosing crops safer than the most profitable one, the second the part closed by combining them so their returns do not all move together. Section 3 maps treatment responses onto the two channels. Observing movements in both measures simultaneously reveals whether an intervention affects crop selection, crop combination strategy, or both.

3 Conceptual Framework

The variance of a crop portfolio is determined by the combination of a household's hunger for risk as well as its ability to diversify away idiosyncratic risk. Failing to account for households' heterogeneous preferences for risk may lead to false conclusions about movements in risk-taking behavior. Two households that both have perfectly efficient portfolios may still have very different portfolio variances resulting from their different preferences for risk, which makes portfolio variance alone insufficient to talk about portfolio efficiency. Without explicitly parsing what role risk preferences play in moving a household's portfolio position in risk-return space, it is hard to discuss diversification and the constraints that prevent households from setting up efficient crop portfolios.

From the empirical application of MPT, I obtain an expected return and standard deviation of return for each crop portfolio observed in a household-year. I am also able to calculate an efficient frontier of crop portfolios, which displays the minimum variance for any selected level of expected return. If a household has a perfectly efficient portfolio, their portfolio property values will lie on this frontier. As a measure of the household's desired level of risk, I use the Sharpe Ratio (expected return/sd of expected return) of the closest point to the household's portfolio on the efficient frontier, illustrated in Figure A.9, as an ideal benchmark. By selecting points on the efficient frontier to represent each household's preferences for risk, I am removing any differential role that constraints may be playing in a household's realized portfolio position. In complement, I expect constraints to be responsible for the magnitude of inefficiency any one household experiences, represented by distance to this ideal closest point on the efficient frontier. I explore credit access,

land ownership as well as wealth as characteristics that I expect to be associated with households setting up more or less efficient portfolios. These two measures, sharpe ratio at the closest point on the efficient frontier and distance to the frontier will serve as outcomes for the exercise in Section 5.5 that asks what role preferences and constraints are playing in portfolio positions.

A household can hold less risk for two different reasons, and the variance decomposition of Section 2.6 separates them: it can choose crops that are individually safer, or it can combine crops whose returns move together less than perfectly. The two ratio measures introduced there track these channels separately in the treatment effect analysis. The selection ratio (Equation 6) captures the self-insurance channel: a treatment-induced increase would indicate that households are shifting land shares toward inherently lower-variance crops relative to the most profitable crop available in their community. The diversification ratio (Equation 5) captures the hedging channel: a treatment-induced increase indicates that the household's crops combine into a portfolio whose risk falls further below the perfectly-correlated benchmark than before, which happens when land moves toward crop pairs that comove less strongly, whether or not any pair comoves negatively. The two measures address different parts of the variance-reduction question: the selection ratio asks whether the chosen crops are intrinsically less risky on average, while the diversification ratio asks whether the chosen crops combine into a portfolio whose total variance falls below the no-diversification benchmark by more than a comparison portfolio's would. Observing treatment effects on both measures simultaneously reveals whether relaxing liquidity constraints affects crop selection, portfolio combination strategy, or both.

Having fixed what is measured, consider what the measured gap can mean. The constraints that push smallholder portfolios inside the efficient frontier are not limited to liquidity. The household-finance literature on the welfare costs of portfolio inefficiency (Calvet et al., 2007; Campbell, 2006) addresses an analogous tension (frictionless mean-variance models that fit the data only after a menu of small frictions is acknowledged) by enumerating the menu rather than structurally modeling any single friction. Following that strategy, the frictions most plausibly relevant to the smallholder agricultural setting examined here are:

(i) **liquidity and working-capital constraints**, where the inability to finance seeds, fertilizer, labor, or land preparation prevents adoption of safer crop combinations (Karlan et al., 2014);

(ii) **subsistence and poverty-trap dynamics**, where households below an asset threshold smooth assets rather than consumption and cling to high-variance cash crops to preserve productive wealth (Zimmerman and Carter, 2003);

(iii) **background risk and incomplete markets**, where the absence of non-farm risk-sharing instruments makes households more conservative on the margin they can control (Heaton and Lucas, 2000);

(iv) **information and learning frictions**, where unfamiliarity with a crop's production or marketing

prevents adoption even when its risk-return profile would favor it (Suri, 2011);

(v) **cognitive and complexity frictions**, where mean-variance optimization across a 16-crop menu is itself a hard problem to solve (Laajaj and Macours, 2024); and

(vi) **cardinality and fixed entry costs**, where indivisible labor, equipment, and management requirements impose a minimum-area floor below which a crop cannot be cultivated profitably (the focus of Section 5.6).

The empirical residual gap from the frictionless frontier is consistent with a combination of all six mechanisms, but unfortunately none of them can be cleanly separated within the panel I observe. The paper's measurement of the gap therefore quantifies the welfare cost of sub-frontier positioning without identifying which friction generates which fraction of that cost; Section 5.6 pursues the cardinality friction in (vi) further because its predictions align most directly with the data.

Two further departures from the canonical MPT setup are handled separately in the paper. *Intensification* (cultivation intensity within a crop is non-divisible, and bimodal cost distributions across input regimes generate effectively different production technologies) is addressed in Section 2.2 via Hartigan's Dip Test, splitting crops with non-unimodal cost distributions into high- and low-intensity variants. *Non-universal crop availability* (not every crop can be cultivated in every community) is addressed in Section 5.3 by constructing crop-set specific efficient frontiers from each community's observed crop set. Appendix I discusses these limitations of the financial-asset analogy in further detail.

The cash transfer's impact on a household's portfolio position is ambiguous a priori. Plausible explanations range from specialization, to intensification, to extensification (scaling up farming operations). If households prioritized specialization, the data would mechanically show a reduction in the number of crops cultivated but depending upon whether or not households were specializing in relatively safer or riskier crops, the portfolio risk could move either direction. The implication would be that the cash transfer provided households with so much liquidity that it overcame the need to diversify within a farming season, and let households pivot to diversifying across time within a given asset class. This is the behavior seen in the agricultural sectors of developed countries. Fully specializing would be a pivot from needing to guarantee income across different states of nature through multicropping, to trusting that in the long run focus on a profitable asset would allow for good years to subsidize bad years. Intensification, an increase in the amount of inputs per hectare, is not a behavior well captured by the scale-invariant nature of MPT. If households were primarily intensifying as a result of the cash transfer, then the MPT framework would show no change in portfolio positions, save for the movement between varieties of crops that are identified as having separate production functions associated within them. Extensification, an expansion of total cultivated area without

altering the relative crop mix, is also scale-invariant under MPT. If households were primarily extensifying as a result of the cash transfer, the framework would register no change in portfolio standard deviation or per-hectare expected return, even as total profit and total hectares both rose in levels. The framework only registers extensification when it is accompanied by a shift in the relative weights w_c assigned to each crop, which is what the extensive-margin entry pattern documented in Section 5.6 captures. These blind spots are why Section 6 reports outcomes in levels alongside the portfolio measures: the levels outcomes register intensification and extensification, and the portfolio measures register what they cannot regarding risk and return.

4 Data & Setting

4.1 The Experiment and the Cash Transfer

The survey data I employ come from a large randomized controlled trial (RCT) in Mali studying household selection into credit market participation, run over the 2010-2012 farming seasons with an additional follow up survey wave in 2017 (Beaman et al., 2023). The sample covers farming households in two Cercles of southern Mali's Sikasso region; the 198 study villages were drawn from a 1998 census of villages meeting size and access criteria within the microfinance institution Soro Yiriwaso's planned expansion zone. The original study's main research question was whether selection into participating in credit markets is associated with an individual's returns to capital. The original authors find strong positive selection into borrowing: individuals who choose to take loans earn high returns on capital (∼130% seasonal returns), while randomly assigned grants generate positive but more modest effects on agricultural output. There were five main treatment arms cross-cutting grants and loans, plus a pure control. I exclusively use the grant-only arm in no-loan villages versus the pure control arm to isolate the effect of the unconditional cash transfer on household portfolio composition. This restriction avoids contamination from the loan treatment arms that drive the original paper's selection results, and ensures that any observed changes in crop portfolios can be attributed to specifically the exogenous cash infusion rather than to credit market access or the interaction between grants and loans. Their experimental design is illustrated in Figure A.10.

The surveyed panel covers 8,632 households across the credit study's 198 villages and the samples of companion experiments, over three years with a long-run follow up five years after the main experiment concluded. The 198 study villages were randomly assigned to loan (88 villages) or no-loan (110 villages) status. For the analysis, I restrict to the balanced panel and follow the original authors' preferred sample, yielding 3,196 households. Critically, all crop return moments (expected returns, variances, and the

covariance matrix) are estimated exclusively from pure control households, ensuring that the estimated crop properties are not contaminated by treatment-induced changes in production behavior. The treatment effect analysis then applies these control-estimated moments to both treatment and control households' observed land allocation shares, which I refer to as weights.

Table B.2 reports balance between the grant and pure-control arms on 55 baseline variables in 8 groups: the Beaman et al. (2023) controls, the main and efficiency outcome suites, crop area shares, livestock, non-crop assets and enterprise, input expenditure, and household demographics. A joint test across all 55 cannot reject balance ($F(55, 3119) = 0.93$, $p = 0.62$), and 5 of the 8 groups are balanced on their own joint tests, including the controls ($p = 0.97$), demographics ($p = 0.38$), livestock ($p = 1.00$), non-crop assets ($p = 0.93$), and input expenditure ($p = 0.11$). Under a Benjamini-Hochberg false-discovery threshold of $q < 0.10$, 3 variables are flagged: baseline portfolio standard deviation ($q = 0.059$), the Sharpe ratio at the closest frontier point ($q = 0.079$), and the diagonal component of portfolio variance ($q = 0.083$), which are three readings of one quantity, since the diagonal is a component of portfolio variance and the frontier Sharpe ratio is a function of σ. The imbalance is therefore concentrated in a single load-bearing dimension rather than spread across the design, and the estimates survive all attempts to interrogate it. Reweighting the sample by inverse propensity scores fitted on all 55 variables drives the standardized difference on baseline portfolio SD from 0.137 to 0.001 and leaves every headline estimate within 0.47 of a standard error of its unweighted value (Appendix G.5). Section 6 takes up the imbalance alongside the empirical strategy.

The cash transfer intervention distributed unconditional grants of 40,000 FCFA (US$140 at the study's 2011 PPP exchange rate of 284 FCFA per USD; approximately US$153 in 2017 dollars) during May and June of 2010, coinciding with the planting season when households face binding liquidity constraints for purchasing agricultural inputs. The grant amount was calibrated to approximate the average loan size of US$113 observed in the study's credit market intervention, enabling researchers to more cleanly separate the effects of capital access from selection into borrowing. Innovations for Poverty Action (IPA) administered the transfers independently from the microcredit program to prevent any association between the grant and loan treatments in recipients' minds.

In villages without access to the loan program, grant recipients were selected randomly. This design feature is particularly valuable for my analysis: it provides exogenous variation in household liquidity at a critical point in the agricultural calendar. The timing of transfers during the planting period means that any observed changes in crop portfolio composition arise from a planting-period cash infusion, ruling out post-harvest channels such as storage or marketing decisions as the operative mechanism. Whether the planting-period response itself is best interpreted as relaxing a capital constraint specifically, as opposed to a wealth effect on risk aversion or an income effect on variety preferences, is taken up in Section 5.6.

4.2 Survey Data and Sample Characteristics

All of the data come from the original study's annually administered main survey, which covered all phases of the farming season. This questionnaire is largely the same across the four years and, importantly, asked about plot-specific crop and factor allocations. This enables the creation of the plot-crop-level profit measure which supports detailed estimation of the moments I am interested in. Throughout, profit per hectare is observed at the farmer-crop-year level, not aggregated to the household and disentangled into per-crop components. Profit per hectare is defined as revenue less expenditure on hired labor, chemicals, land preparation and seeds, normalized by the number of hectares cultivated for production. The survey records revenues at the plot-crop level and input expenditures (chemicals, land preparation, labor, seeds) at the plot level, so on sole-cropped plots profit per hectare maps directly to the crop; on intercropped plots, plot-level inputs are allocated across all the constituent crops in proportion to plot revenue shares (Section 2.2). The data therefore permit direct estimation of crop-specific moments from observed plot inputs and outputs, without an additional disaggregation step.

On average, households in the sample are cultivating 5.7 unique types of crops (SD=2.2) as part of their portfolio and the mode of the distribution of the number of crops is 5, as shown in Figure A.11. These crops are planted across 5.00 plots (SD=2.7) which total 7.0 (SD=4.5) cultivated hectares on average. Appendix Table B.1 reports cultivation frequencies by season, and Table B.5 gives a one-line description of each crop. Profit per hectare averages \$306 2017 USD (SD=297). Multiplying the profit per hectare against the number of cultivated hectares yields an average total profit of around \$2198 per year per household which, once divided by typical household size, places households without other income-generating activities firmly below the World Bank \$2.15 (2017 USD) per person per day poverty line.

Chemical expenditures on inputs like fertilizer, weedicide, pesticide and fungicide constitute the largest average input for households, averaging \$842 (SD=996). Labor expenditures average \$67 (SD=132), while expenditures on seeds average \$32 (SD=58). Land preparation expenditures on ploughing services average \$33 (SD=68). The high standard deviations on these averages showcase the riskiness of smallholder agriculture as well as the varying levels of investment across individuals. Tables B.3 and B.4 provide the full distributional summary for all survey-measured and MPT-constructed outcomes, respectively.

4.3 Production Risk in Southern Mali

Commonly discussed consumption-smoothing tools such as community risk sharing are less effective in this setting because households within a community experience common shocks through the aforementioned weather and price risk channels discussed in Section 2.3. Appendix G.3 puts a number on this: the median

household's idiosyncratic share of portfolio variance is only 15.2%. The risk a household bears is therefore mostly shared with its neighbors rather than specific to it, which is what makes crop choice, rather than risk pooling, the margin worth studying. The region has few year-round rivers providing irrigation opportunities for farmers. While it would theoretically be possible on a national level to improve irrigation opportunities, the low population density of this area makes it difficult for the Malian government to justify massive public works projects of the kind seen with canal expansions in India and Southeast Asia. Rain is the primary way smallholder farmers have to get water to their crops and this makes weather risk a priority to understand.

In Mali, there is one rain-fed farming season per year (Beaman et al., 2023; FEWS NET, 2024): the rains fall from May to about September, and the farming year runs from sowing in May to the last harvests in December, with January to April dry. Cultivation in the dry season is rare. The farming season itself has three parts: sowing from May into July as the early rains arrive, growing-season rainfall through September, and harvest from September through December depending upon the crop. This common shock structure is also why the portfolio benchmarks constructed later in the paper are built from locally available crop sets rather than a single national menu (Section 5.3): the relevant risk-return menu is defined by local conditions, not a national market.

4.4 Baseline Wealth Index for Heterogeneity Analysis

I construct a baseline wealth index as an equal-weight standardized sum of nine pre-determined asset stocks at baseline: per-capita counts of productive durables (plough, cart, motorbike) and consumption durables (radio, stove, television), the household's livestock head count, and two monetary asset-value variables (value of owned goods and household business-asset value). All nine components are predetermined with respect to treatment; cultivated area and labor inputs (choice variables) are deliberately excluded. The seven count components enter as z-scores; the two monetary components are right-skewed with mass near zero and are $\log(1+x)$-transformed before standardization, which preserves discrimination among poor households. Missing components are median-imputed within variable; one livestock-value variable with zero sample variance is dropped.

5 Constructing Crop Portfolios

5.1 Building up a covariance matrix

After implementing the methodology detailed above, an underlying positive relationship between estimated risk and return can be seen across crop types, illustrated in Figure A.8. There is a cross-crop Pearson

correlation between expected return and standard deviation of 0.95 ($\rho_{Spearman} = 0.93$). Higher-return crops are also the riskier ones, which is the tradeoff a household must navigate when it chooses how to allocate its land. If risk and return were unrelated across crops, a household could raise its expected return without taking on more variance, and there would be no portfolio problem worth studying. That they are related this tightly is not a coincidence, and the two sides of the relationship have different origins. Conditional on price, a crop's variance is a property of its biology and of the weather: Appendix C.10 decomposes each crop's within-farmer variance of log revenue per hectare into yield, price and covariance terms, and yield variation accounts for 0.98 of it across the 14 crops with at least 100 farmer-crop cells (0.92 to 1.00 crop by crop), price variation for 0.04, and the covariance, negative for the staples, for -0.02; restricting to harvests priced at their own community's median leaves the shares at 0.98, 0.04 and -0.02. Expected return, by contrast, runs through the price, which is set where the supply decisions of risk-averse cultivators meet demand. A crop whose variance is fixed by nature survives on the menu only if its price compensates the farmers who bear that variance, so in equilibrium the cross-crop line in Figure A.8 is the premium these households require to hold risk, and the near-common Sharpe ratio it implies is an equilibrium condition rather than a convenient regularity. That reading is also what gives the diversification ratio of Section 2.6.1 its welfare content: on a menu where returns are proportional to volatility, maximizing it is maximizing the Sharpe ratio. The paper takes the menu of moments as given; a $153 grant to a third of the households in half the villages, in a setting where 74% of harvests are never sold, is not a shock that moves crop prices.

A monocropping household's portfolio position on Figure A.8 would overlap exactly with its chosen crop's point in this space. The crops' exact positions are detailed in Table B.6, alongside the skewness and excess kurtosis of the same farmer-demeaned residuals used to create the first two moments. 90.9% of household-years in the sample simultaneously cultivate crops from both ends of this risk-return spectrum, illustrating that households often hold combinations spanning safe and risky assets rather than specializing at one extreme. Fonio millet is the crop providing both the highest expected return and the highest riskiness of the 16 crop types I consider, with cotton sitting just below it on both dimensions. Maize, one of the most commonly cultivated crops in the region, is a single crop on the survey roster but I treat it as two assets by splitting its plots at the control-group mean input cost per hectare ($210/HA). This decision is informed by a Hartigan's dip test that shows evidence of multiple production functions for maize being present (Section 2.1.3). The high-input variant [Maize (High)] sits near the middle of the risk-return space and the low-input variant [Maize (Low)] exists at noticeably lower risk and lower expected return values.

Table B.6 also reports skewness and excess kurtosis: every crop is right-skewed and heavier-tailed than a normal distribution, which is what a setting with near-zero harvests on the downside and occasional price spikes on the upside produces. The analysis works with the first two moments throughout; Appendix G.4

checks that the grant does not move the portfolio-weighted higher moments.

The estimated correlation coefficient terms, before being scaled to create a variance-covariance matrix, are shown in Table B.7. Table B.8 reports how many observations stand behind each entry, and the counts vary widely: the median crop pair rests on 474 farmer-years, but 17 of the 120 pairs rest on fewer than one hundred, and the thinnest, Bambara groundnut with sesame, on 20. The diagonal draws on a different and larger sample, since a crop's variance uses every plot-crop observation rather than the overlap of two cultivator sets, and runs from 366 to 10,831 observations. Only two of the 120 unique terms estimated are negative, sitting below the $\rho = 0$ independence threshold as true hedging opportunities. This reflects the fact that crops share factors of production: a temperature ceiling above which water evaporates too quickly to be retained, some minimum amount of rain, etc. Figure A.12 shows the distribution of these correlation coefficients and highlights where the opportunities for diversification live: not in negative correlations, but in the spread of the positive values. The pairwise correlations run from -0.18 to 0.56, with a median of 0.30 and an interquartile range spanning 0.22 and 0.38.

The blue line in Figure A.13 represents the efficient frontier for the full crop-set after accounting for the crops' covariances and marks the boundary of what portfolio positions are feasible. The reductions in variance available along the blue line at any given level of return result from the presence of diversification opportunities in profits across crop types. The variance-covariance matrix behind that line, Table B.9, is the correlation matrix of Table B.7 scaled by the crop standard deviations reported in Table B.6. To obtain an expected return as high as one obtained by monocropping fonio, a household cannot cultivate any other crops, and must accept the variance that comes along with monocropping this high-risk crop. This is why the efficient frontier leads directly to the fonio monocropping point. At any lower level of desired return though, households can reduce the risk of their portfolio by cultivating different crops in combination. For example, if a household desired the expected return of cotton, the second-highest return crop I observe, instead of monocropping cotton and taking on its variance, they can plant a combination of crops that form the same expected return but substantially lower variance by moving along the efficient frontier.

Several crops are dominated, offering no more expected return at no less risk than another crop on the menu, yet they are almost never the sole crop a household grows, which is what their value through the off-diagonal correlations implies (Appendix G.12).

5.2 Where Households Sit Relative to the Frontier

While Figure A.13 highlighted the ingredients required to estimate the expected return and standard deviation of a portfolio; Figure A.14 adds the actual positions of each observed household-year portfolio to

this space as green dots. The only thing driving variation in the position of these observed portfolios is a different shares of land being allocated across crop types. It is mechanical that that all the observed portfolio positions fall beneath the efficient frontier, but it was not guaranteed that the mass of observed portfolios would lie closer to the efficient frontier than the monocropping positions. This pattern is suggestive evidence that the MPT machinery is picking up the risk-management benefits of multicropping. The spread of dots against the efficient frontier reflect households' different preferences for risk, while the variation in distances from the efficient frontier represent how impacted by constraints a household is (Section 3). Pooling the four survey rounds, the typical pure-control household holds a portfolio with an expected return of $325 per hectare and a portfolio standard deviation of $342 per hectare, so the risk carried slightly exceeds the return expected. Variation across households is small against that risk: expected return runs from $276 to $375 per hectare across the interquartile range (Table B.4).

Comparing these estimated expected portfolio return values against realized profit per hectare values in Appendix Table B.10 showcases a strong positive relationship across all models. The coefficient on the constructed expected return is statistically significant in every column, with point estimates ranging from 0.75 to 1.56 and my preferred coefficient sitting at 0.80 with year fixed effects, the specification the treatment-effect analysis uses.[11] Each control household contributes its own farmer-crop means to the $\bar{R}_c$ it is scored against; Appendix C.9 removes that overlap with a leave-one-out $\bar{R}_c$ exercise and the preferred coefficient moves from 0.802 to 0.798, so the circularity is immaterial.

5.3 Defining the Relevant Frontier for Each Household

Not every crop is cultivated in every community, and in those cases the full-menu frontier of Figure A.14 is the wrong benchmark to use when evaluating efficiency. For example, a community that never grows fonio cannot reach the top of the full-menu frontier. I therefore treat crops that are never cultivated in a community in any of the four survey rounds as infeasible to produce there, and remove them from that community's covariance matrix. I then recompute efficient frontiers for each community based on only the crops that were grown.[12] I call these crop-set specific rather than community-specific frontiers, since the frontier a household faces is set by its menu, and community-level cultivation is only how I observe the menu. The 110 communities in my sample reduce to 18 distinct crop sets and therefore 18 distinct frontiers. Exclusions are the exception rather than the rule: the median community cultivates 15 of the 16 crop types at

[11] The R^2 values of these regressions are 0.109 without fixed effects, 0.469 with year fixed effects only, 0.115 with crop-set fixed effects only, and 0.469 when both are included. Appendix C.11 outlines what the R^2 pattern does and does not say about the construct.

[12] The data give no guidance on whether a crop is absent because no trader who visits that community buys it, because local conditions cannot support it, or for some other reason. Appendix G.8 discusses the candidate reasons and shows the results are robust to how the crop set is defined, so I treat what is never grown as unavailable.

some point, only 8.3% of community-crop pairs are never cultivated, and 31% of communities (holding 43% of household-years) display the full set (Appendix Table B.11). Figure A.15 displays the resulting frontiers. Removing a high-return crop such as fonio or yam shortens the frontier at the top; removing a crop that hedges the rest shifts it rightward and flattens it, by an amount that is a function of its covariances with what remains. Because an exclusion can only shrink the attainable set while the household's own portfolio stays feasible, the benchmark moves toward the household and measured inefficiency falls mechanically. The two frontier-based outcomes used throughout the paper, the Sharpe ratio at the closest frontier point and the Euclidean distance to that point, are computed against these crop-set specific frontiers.

5.4 How a Portfolio Approach Compares with Standard Diversification Measures

To understand the heterogeneity that the portfolio variance measure captures and that the number of crops and the Normalized Herfindahl-Hirschman Index (NHHI) miss, Figure A.22 plots these three dimensions together. I limit the sample to portfolios containing just low intensity Maize, Rice or combinations of the two crops to keep the geometry legible. The X axis of this figure shows the SD of portfolio return, while the Y axis maps the NHHI value. The two shapes show the number of crops included in the portfolio: squares for monocropped portfolios and circles for 2-crop portfolios and the color maps to a portfolio being either mostly maize or mostly rice. All monocropped portfolios are mechanically found at the very top of the figure, where NHHI = 100. For these households the number of crops and the NHHI carry the same information: both sit at their boundary values (one crop, an NHHI of 100) and correctly show perfect specialization. Both measures, however, miss the heterogeneity in riskiness across crop types that the portfolio variance measure captures. Monocropping maize is less risky than monocropping rice and treating those farms equally can lead to false conclusions about risk taking behavior.

Below the corner solutions of perfect specialization, the NHHI clearly improves on the crop count: all of the variation along the Y axis beneath NHHI = 100 is invisible to the number of crops cultivated as an outcome variable. Moving up the Y axis reflects increasing specialization, with more hectares dedicated to one of the two crops. The NHHI still does not differentiate between the riskiness of crop types, however, and that failure is what keeps a specialization measure from capturing the household's risk-taking behavior. The parabola contained in the figure traces out households allocating different shares of their land between the two crop types. Both the X axis and Y axis positions of portfolios are being driven by differential hectare weights, but only on the X axis are the hectare weights being scaled by a framework that accounts for how

correlated risk is across crops.[13] This lets the portfolio variance measure capture much more information than the NHHI. Using this example, households with identical NHHI values of 50 could have a return SD of around either \$325 or \$500, a sizeable difference. I highlight this by including the number-line on the right of the main figure. When just looking through the NHHI, the parabola collapses on itself and the mostly-maize and mostly-rice portfolios end up side-by-side even though my estimated portfolio risk measure highlights how different these portfolios truly are. When using NHHI as an outcome, we fail to understand if households are specializing in safer or riskier activities and are just asking if they're more specialized at all.

Figure A.23 generalizes the illustration that variation is hidden by N Crops and NHHI outcome measures to the full crop set, but looking at a narrower sample: just households holding portfolios of six crops and all with NHHI values $\in [15, 20]$. Households holding the same number of crops in the same proportions can bear risk differing by a factor of 1.6. Among the 602 household-years cultivating exactly 6 crops with an NHHI between 15 and 20, indistinguishable on both standard measures, portfolio standard deviation runs from \$262 at the tenth percentile to \$426 at the ninetieth. Again, these proxy measures record how land is divided, not how risky the crops being divided over are.

However, Figure A.23's cross-sectional comparison still pools many sources of variation. A sharper test of what the standard specialization measures miss looks within-farmer at the marginal effect of adding a single crop. For each household-year pair where households added exactly one crop and dropped none between consecutive waves (1,505 such transitions in the panel), I compute the change in portfolio standard deviation and group by which crop was added. A household that shifts between the intensity variants of one crop has substituted within that crop rather than added one, so those shifts are netted out of both the added and the dropped set before the restriction is applied. Figure A.24 reports the mean $\Delta\sigma$ on portfolio SD by added crop. Adding cotton raises portfolio SD by \$31.6 per hectare ($n = 112$); adding fonio raises it by \$19.0 per hectare ($n = 31$); adding most other crops (okra, sorghum, low-intensity maize, beans, sweet potato, sesame, the "other" aggregate) lowers portfolio SD by between \$10 and \$38 per hectare on average. Two households that both register $\Delta N_{\text{crops}} = +1$ between waves can therefore move in opposite directions in σ-space, with a wedge of roughly \$56 per hectare separating the cotton-adder from the okra-adder. The number of crops cultivated outcome would assign these two households the same value of the diversification measure; the MPT framework distinguishes them by the magnitude and sign of the variance contribution that the added crop makes through its diagonal σ_c and its off-diagonal correlations with the household's existing portfolio. This is the empirical content of the measurement contribution: the standard specialization metrics capture the entry of a new crop but not its risk-relevant identity.

[13] The shape of the parabola is in part determined by the relative riskiness of the two crops as well as the correlation coefficient between them. In many cases, the lowest risk portfolio will not be the left-peak of the U, but along the middle of the left side of the U shape.

5.5 Endowments and Portfolio Efficiency

Section 5.2 left every household inside its frontier, and the gap is not small: the average control household sits \$49 per hectare below the attainable frontier of Section 5.3 at its own level of risk. Preferences for risk should not be the answer: in the framework of Section 3, they move a household along its frontier, not away from it, so what remains is a set of binding constraints. This subsection asks which household endowments are associated with the distance between a household's portfolio and an efficient mix of its crops. I focus on observable household characteristics that theory suggests should predict portfolio efficiency: asset holdings, land size, and shock exposure.[14]

Figure A.16 presents standardized bivariate associations between baseline household characteristics and portfolio efficiency, where efficiency is defined as proximity to the crop-set specific efficient frontier (higher = closer and more efficient). Each row reports the coefficient from a separate bivariate regression of z(Efficiency) on z(predictor), on the baseline (2010) sample, with crop-set fixed effects and standard errors clustered at the household level (3,969 households, 18 crop sets, 110 communities). These regressions are correlational since the covariates are not randomly assigned, so the associations should be interpreted as descriptive patterns rather than causal relationships. The pattern in Figure A.16 aligns with the constraint-based interpretation of inside-the-frontier positioning: physical capital (plough use, fertilizer use, livestock value), farm size, household size, and financial savings all significantly predict greater efficiency, while "not enough food for the household" significantly predicts lower efficiency.

Entering the predictors jointly (Table B.12; Appendix G.13) keeps plough use, household size, farm size and financial savings and drops fertilizer use, livestock value and food insecurity: what survives describes a household's capacity to act, through traction, land, labor and cash on hand, rather than its accumulated stock of goods.

Figure A.17 reports the parallel exercise for the risk dimension, using the level of risk a household appears to aim at rather than the level it realizes. A household's observed σ mixes the two margins the framework separates: preferences move it through the choice of target, while constraints move it through how far the household sits from that target. I measure the target as σ at the point on the household's own crop-set specific frontier closest to its observed portfolio, so that the target and the distance in Figure A.16 decompose the observed position between them. Every household in the sample sits to the right of its own frontier, so this projection always lowers σ relative to the realized value. Three capital measures predict a higher target: plough use ($\hat{\beta} = 0.104$, $p < 0.001$), fertilizer use (0.060, $p < 0.001$), and livestock value (0.045, $p = 0.003$),

[14] The survey elicited risk preferences from hypothetical lottery menus, with each household randomly assigned to a small-stakes or a large-stakes version and re-asked across the three main survey rounds. I do not use these elicitations anywhere in the paper, just as Beaman et al. (2023) makes no use of them in their published analysis. This is also why later on in the welfare bounds in Section 7 I work with implied risk-aversion ranges rather than elicited parameters.

while "not enough food for the household" predicts a lower one (-0.090, $p < 0.001$). Farm size moves the target far less than it moves proximity to the frontier (0.036 against 0.127), which is the pattern a limit on the number of crops would produce, the mechanism Section 5.6 takes up next: more land supports a wider crop set and a position closer to the frontier, rather than a different level of risk.

The conceptual framework separates constraint variables, which should predict distance from the frontier, from preference variables, which should predict the level of risk a household targets. Financial savings and household size are among the strongest predictors of efficiency, yet neither predicts the risk target ($p = 0.985$ and $p = 0.620$ respectively): they move households toward the frontier without moving where on it they aim, which is the pure-constraint pattern the framework predicts. "Not enough food for the household" does not fit that separation: it predicts lower efficiency in Figure A.16 (a constraint blocking the frontier) and a lower risk target in Figure A.17 (defensive crop choices), which is the joint pattern a binding survival constraint produces. One caveat bounds how far the preference side can be pushed. The target is derived from the household's own observed portfolio rather than measured independently, so it is the realized position purged of the displacement rather than a recovered preference parameter, and I read these associations as descriptive. Read together, the two figures say that what predicts proximity to the frontier is a household's capacity to act at planting rather than its appetite for risk.

5.6 Fixed Costs Limit the Number of Crops a Household Grows

Figure A.18 shows why the number of crops is an important channel to explain: mean portfolio risk falls from \$436 per hectare among one-crop household-years to \$307 among those growing eight or more, a 30% decline at nearly unchanged expected return, so a household that can grow more crops has more opportunities to produce an efficient portfolio. This subsection presents the evidence for a per-crop fixed cost on cultivation and then builds the portfolio problem such a cost implies. The cultivation data come first: a per-crop threshold predicts that crop count rises with farm size and with wealth, and that the gradient flattens among the smallest farms, since a household below the threshold cannot convert additional land into an additional crop. The input-cost data come second, where the same threshold predicts that spending scales with the number of crops in the categories where indivisibility bites and with area everywhere else. The resulting cardinality constraint is then written into the mean-variance problem and its frontier characterized.

Two further patterns discriminate the discrete reading from the continuous one but are reported where they belong rather than here: the treatment-effect null on extensive-margin cultivation among the smallest-farm quintile, in Section 6, and the identification diagnostics that weigh the entry-cost family against soil heterogeneity and continuous capital constraints, in Appendix G.6. None of these patterns is predicted by a

smooth continuous-wealth mechanism without additional structure.

5.6.1 Larger and Wealthier Farms Cultivate More Unique Crop Types

Doubling cultivated area is associated with cultivating approximately one additional crop type ($\hat{\beta}_{\log K} = 1.138$, $p < 0.001$, estimated on raw self-reported cultivated area values with community fixed effects; Appendix G.6 explains the 0.2 HA floor convention and why the gradient is estimated on pre-floor area). Mean cultivated cardinality rises monotonically with farm size, from 3.0 crops for sub-2-HA farms to 7.0 crops for farms above 10 HA (Figure A.19). This is the size-diversification gradient a per-crop fixed cost predicts. If smallholder cultivation involves a minimum-area threshold $L_{c,\min}$ below which cultivating a crop is not profitable, generated by indivisible labor blocks, plough preparation requirements, seed packets, and supervision time, then the maximum number of crops a household can cultivate is bounded above by $\lfloor K/L_{c,\min} \rfloor$, where K is the household's total cultivated area: larger farms can support more crop types, while smaller farms hit the cardinality bound earlier. Clearing that threshold also takes cash up front, for the block's seed, ploughing and labor, an outlay I denote F and estimate in Appendix G.7; a household needs both the land and the cash to add a crop, so a grant operates on F while $L_{c,\min}$ and K bound how many crops that cash could ever fund.

The same fingerprint appears when households are ranked by baseline wealth rather than farm size. The richest tercile according to a wealth index is 41% more likely than the poorest to cultivate at least six unique crop types and 94% more likely to cultivate at least eight, and the top tercile's crop-count distribution first-order stochastically dominates the bottom's via a one-sided Kolmogorov–Smirnov test ($p < 10^{-9}$). First-order stochastic dominance in cultivated cardinality is what a per-crop fixed cost predicts once wealth relaxes the resource that funds entry; it is not what a pure risk-preference story predicts, since households that merely tolerate more risk would concentrate in a few high-return crops rather than hold systematically more of them. This suggests that there is a credit constraint on the number of crops a household can cultivate at once. The next subsection measures what the missing cash has to buy.

5.6.2 Input Costs Scale with the Number of Crops, Not Area

The cost side of the input data carries the signature most directly. Within farmer, holding cultivated area fixed, the input categories that indivisibility predicts to be lumpy do scale with crop count rather than with area (Table G.11). Seed expenditure rises by \$4.28 per additional crop and plough and plot preparation by \$4.04, both significant at the 1% level with farmer and round fixed effects. Fertilizer is the contrast that identifies the mechanism: the largest category at 61% of spending, it tracks cultivated area and is indistinguishable from zero on the crop-count margin ($p = 0.24$). These are the same indivisibilities that

Foster and Rosenzweig (2022) show make the smallest farms costly to operate per hectare, machine capacity and the fixed cost of transacting for labor, binding here one level down at the crop margin rather than at the farm. The shape of total input costs points the same way. The null that each additional crop carries an identical expenditure increment is rejected both when every cultivated crop is counted ($\chi^2(7) = 94.5$) and when the count is restricted to crops holding more than 5% of the household's land ($\chi^2(7) = 39.8$), both $p < 0.001$. The second panel matters because it answers the obvious objection: the pattern is not an artifact of households adding token plantings that cost almost nothing, since it survives when only crops holding real land are counted (4.20 per household on average, against 5.5 crops of any size). Among those land-holding crops the rejection is driven by a lump at entry: moving from one crop to two raises input expenditure by roughly \$272, while no increment from two through eight crops exceeds \$61 in absolute value (the final step to nine crops sits in a thin cell and is imprecise).[15] Later increments shrink mechanically with the marginal crop's land share, since much of measured spending varies with area, so the declining tail speaks to allocation rather than to the fixed cost fading. The lump at entry and the per-crop loading of seeds and land preparation are the components no smooth cost structure produces. A continuous-friction alternative remains possible: a wealth-dependent capital constraint with a concave mapping from wealth to frontier proximity (e.g., diminishing returns to capital in the access-to-inputs technology) would also leave a residual gap when continuous wealth measures are entered linearly, as would measurement error in those wealth proxies or omitted preference heterogeneity. A residual gap alone does not therefore identify a discrete friction. What discriminates the discrete from the continuous interpretation is the empirical signature of indivisibility assembled above: a discontinuity in cardinality and weight allocation at small farm sizes that would be smoothed away under continuous mechanisms. Appendix G.6 re-estimates the seed and land-preparation loadings controlling for the number of plots, which addresses the possibility that a crop added by intercropping needs seeds but not fresh plot preparation. The diagnostics tighten identification toward the entry-constraint family relative to a generic "inside-the-frontier" observation, but they do not isolate the per-crop fixed cost as the unique mechanism; labor-supply or experience-based versions of the same threshold story could produce qualitatively similar patterns. The discrete-vs-continuous distinction in this paper is therefore a directional claim, not a structural identification.

[15] Increments are within-farmer differences of crop-count level coefficients from a regression of winsorized input expenditure on crop-count indicators, cultivated area, and a cotton indicator with farmer and round fixed effects; equality is tested with a Wald statistic on the full increment vector. These regressions use household totals of plot-level expenditure records, so the revenue-share rule that allocates plot inputs across crops for crop-specific profitability (Section 2.2) never enters and the contrast is immune to that construction.

5.6.3 A Cardinality-Constrained Portfolio Problem

The empirical pattern motivates a cardinality-constrained extension of the mean-variance optimization in Section 2. Let w_c denote a household's share of cultivated area allocated to crop c. The frictionless problem maximizes $U = \mu - \frac{\phi}{2}\sigma^2$ subject only to $\sum_c w_c = 1$. A fixed-cost extension imposes a semi-continuous constraint: for each crop c, either $w_c K \geq L_{c,\min}$ (the household commits at least $L_{c,\min}$ hectares) or $w_c = 0$. This formulation has been studied extensively in the finance literature on portfolio selection with discrete asset-choice constraints (Chang et al., 2000; Bertsimas and Shioda, 2009; Jobst et al., 2001); it is the natural finite-resource version of the mean-variance farm-planning models that motivated Freund (1956) and the MOTAD literature initiated by Hazell (1971), integrated into a broader risk-programming framework in Hazell and Norton (1986). This paper revisits that intellectual arc with modern panel data and causal identification, while remaining agnostic about the exact magnitude of the per-crop entry threshold.

5.6.4 The Constrained Efficient Frontier

Figure A.20 plots constrained efficient frontiers against the unconstrained frontier from Section 5. The constrained frontier is constructed as the upper envelope of the $\sum_{k=1}^{8} \binom{16}{k} = 39{,}202$ separately-optimized restricted frontiers, one per feasible subset of ≤ 8 crops. The figure therefore illustrates how a cap on the number of crops reshapes the attainable set of portfolio positions, and should not be read as a calibrated model of the constraint households actually face.

Two features of the constrained frontier deserve emphasis. First, the constrained frontier shifts inward across the entire σ-range, but the shift is concentrated at the low-σ end, where extreme diversification is no longer feasible. The constraint binds hardest precisely on the safety-seeking households the paper's treatment effects target. Second, the constrained frontier characterizes the achievable set facing households who cannot afford the fixed cost of having a seventh or eighth crop in their portfolio. Given constraints, efficiency could be measured against this constrained envelope, not the frictionless one. Section 5.5 reports distance to the unconstrained frontier; under the constrained interpretation, the same households are weakly closer to feasibility, and the empirical content of the residual gap is the welfare cost of binding fixed costs rather than systematic optimization failure.

Figure A.21 restates the same cardinality constraint as a curve rather than as a family of frontiers. At each maximum crop count K, the Gini coefficient measures how far the K-crop frontier bows away from the straight chord joining the two corner crops, so it tracks the diversification room a K-crop menu opens up rather than the number of portfolios it contains. That room is bought almost entirely in the first six crops. The coefficient rises from 0.354 at two crops to 0.469 with the full menu, and 90% of that rise is in place by

6 crops; beyond that point the marginal return to another crop is small.

Against that curve Figure A.21 also overlays the distribution of crops households actually cultivate, split by baseline wealth tercile. The first-order stochastic dominance reported earlier in this section is visible directly: mean crops cultivated rise from 5.4 in the poorest tercile to 5.5 in the middle and 5.8 in the richest. What the curve adds is where along it that difference falls. The richest tercile's mean sits at the elbow (5.8 crops against an elbow at 6), on the flat segment where an additional crop buys little further diversification room; the poorest tercile is the largest of the three at every count from 2 to 4 crops, on the steep segment where each crop forgone costs materially more. Wealth therefore does not simply buy more crops, it buys access to diversification opportunities on the part of the curve where crops are worth the most.

Section 5 has only presented associations, and it leaves two open questions that only an exogenous shock to liquidity can settle. The first is direction: if cash relaxes a survival constraint of the kind the food-insecurity result suggests, recipients should take on more risk; if it clears the entry cost of Section 5.6, they should add crops and take on less. The second is the margin: whether cash closes the gap by adding crops or by improving the mix among crops a household already grows. Section 6 answers both with the grant.

6 The Cash Transfer's Effect on Crop Portfolios

6.1 Empirical Strategy

To understand the impacts of the cash transfer on the expected return and variance of each cultivator's portfolio, I estimate the analysis-of-covariance (ANCOVA) model in Equation 7. My regressor of interest is an indicator for cultivator i receiving the 40,000 FCFA cash grant in a post-treatment year $t \in \{2011, 2012\}$, D_{it}. Due to its random assignment, I can interpret its effects on the outcomes I examine as causal. I employ ANCOVA specifications because including the baseline value of the outcome variable increases the specification's power (McKenzie, 2012). I also include the preferred control suite of Beaman et al. (2023), ρ_{it}, which aids in precision, and a year fixed effect γ_t (an indicator for 2012, as there are only two post-treatment rounds) to account for common shocks across the pooled sample.[16] Standard errors are clustered at the household level.

$$Y_{it} = \alpha + \beta_1 D_{it} + \beta_2 Y_{i,t=2010} + \gamma_t + \rho_{it} + \epsilon_{it} \tag{7}$$

Alongside my novel measures of portfolio risk, I estimate the model on a standard set of ex-ante planting decisions and ex-post realized agricultural outcomes. Because the MPT outcomes are constructed from

[16] Unlike the original specification, I do not include community fixed effects: the original authors absorb village fixed effects to account for their village-level loan randomization, but since I restrict to arms 4 and 5 (both in no-loan villages), this stratification dimension is not relevant to my comparison.

estimated moments, I supplement the analytical inference with randomization inference (RI) using 10,000 permuted treatment assignments from the original experimental design. Because I test treatment effects on 16 outcome variables simultaneously, I also apply the Westfall-Young stepdown procedure (Westfall and Young, 1993) on those same permutations to control the familywise error rate within four outcome families: (i) ex-ante cultivation decisions, (ii) ex-post realized outcomes, (iii) MPT-constructed portfolio measures, and (iv) efficiency measures. The Westfall-Young adjusted p-values are reported as a row in the results tables and further details are provided in Appendix E.

Table B.2 reports balance on the portfolio outcomes alongside the household characteristics. While most outcomes are well-balanced, baseline expected profit per hectare and portfolio standard deviation show statistically significant differences. The magnitude of the baseline imbalance on Portfolio SD (+9.16 USD/HA, more than twice the absolute size of the headline treatment effect) warrants four additional diagnostics: a randomization-inference test of the imbalance against the original BKTU permutation distribution; a difference-in-differences specification that absorbs baseline levels rather than partialling them out; a specification that adds direct controls for the three imbalanced crop shares; and a reweighting of the sample by inverse propensity scores fitted on the 55 baseline variables of Table B.2. All four diagnostics are reported in Appendix G.5. Reweighting drives the standardized difference on baseline portfolio SD from 0.137 to 0.001 but only moves the portfolio-SD estimate from -4.01 to -3.80 USD/HA ($p = 0.033$); it remains significant and no headline estimate moves by more than 0.47 of its standard error. The randomization-inference test confirms the imbalance is statistically distinct from a typical permutation draw (RI two-sided $p = 0.0294$ across $10,000$ BKTU permutations). The difference-in-differences specification yields a larger effect in magnitude (-10.558 USD/HA, $p = 0.0002$), but the gap is what differencing a chance imbalance produces: with baseline and follow-up portfolio SD autocorrelated at $\rho = 0.30$, DiD subtracts the whole baseline gap where only ρ of it persists, and the wedge it opens (-6.55 USD/HA) is close to what that overcorrection predicts (-6.37). The DiD column is reported to show how much work the baseline adjustment does, not as an alternative estimate.

Finally the crop-share-controlled specification yields a treatment effect of -4.273 USD/HA, essentially identical to the headline result. Given these results I conclude that the imbalance does not threaten the headline conclusion from the ANCOVA analysis.

6.2 Average Treatment Effects

6.2.1 Headline Effects on Standard and Portfolio Outcomes

The results of this analysis are presented in Table B.13. The cash transfer has statistically significant effects across nearly all outcomes, painting a coherent picture of how liquidity relaxation affects agricultural portfolios. Beneath the estimated standard errors are randomization inference p-values from an exercise using 10,000 potential treatment assignments. The stars in the table correspond to the significance levels determined by the standard errors. The two methods for classifying significance only disagree on one coefficient and only differ on whether or not the coefficient is significant at the 5% or 1% level.

Panel A reports standard outcomes for agricultural-focused RCTs. The grant increases total expenditures by \$31, total revenue by \$80, and profit in levels increased by \$74.[17] Households invest only about one fifth of the total transfer into agricultural inputs but realize positive returns on those investments in levels. Households also scale up their production, cultivating a little over a third of an additional hectare after receiving the grant. After accounting for this increase in cultivated area, profit per hectare shows a positive but imprecise effect of \$4.49 (SE=3.16). It is important to highlight that expenditure per hectare is unchanged (\$0.59, $p = 0.69$)[18] This indicates that the increase in total spending operates entirely through the extensification margin (more land cultivated in the same way) rather than through intensification of inputs on existing land discussed in Section 3. If intensification was the primary response to receiving the grant, this would have left MPT-informed portfolio positions unchanged. Instead, Panel B illustrates that households use the liquidity to expand and diversify.

Panel B of Table B.13 highlights old and new outcome measures focused on evaluating risk-taking. The extensive-margin response documented crop-by-crop in Section 5.6 aggregates at the household level to a 0.33-crop increase in portfolio breadth alongside the 0.36-hectare expansion in cultivated area. Because portfolio size is impacted by the grant, it is inappropriate to use the standard Herfindahl-Hirschman Index (HHI)[19] and I use the Normalized Herfindahl-Hirschman Index (NHHI) instead. The NHHI falls by 0.87 points ($p < 0.05$), reflecting the same compositional shift in concentration terms; even after accounting for the additional crop types present, land is distributed more evenly across crops. A traditional reading of these magnitudes would conclude that the grant relaxes constraints and enables broader diversification, reducing risk and expected return together as a bundle. Neither outcome can reveal whether risk and return are moving independently. The limitation is quantitative, not rhetorical. Whether this spreading land more

[17] Since I am using an ANCOVA specification, the cost and revenue coefficients need not linearly combine to become the profit coefficient.

[18] That bounds intensification on the input side while Appendix C.8 bounds it on the outcome side, asking whether the control-estimated moments price treated portfolios correctly and the moments appear accurate.

[19] The lower bound of the HHI value for a household is 1/(Number of Crops), a function of its portfolio size.

evenly across more crop types lowers risk or not depends on which crops absorb the land and how their returns covary. The portfolio-risk consequences in the next subsection of Panel B supply exactly that.

Focusing on the MPT-informed outcomes in Panel B shows a much higher-resolution story about households' response to the grant. The grant reduces portfolio standard deviation, risk, by \$4.01 per hectare ($p < 0.05$), a 1.3% reduction relative to the control mean of \$316. Expected returns per hectare show no significant change (−\$2.38, $p = 0.14$), suggesting households are achieving risk reduction without sacrificing expected returns. The elasticity calculated from these results shows that a 1% increase in agricultural expenditure reduces portfolio risk by 0.23%. The grants were used to buy safety and efficiency, not higher returns. This movement is welfare-improving: the Sharpe ratio increases by 0.67 points ($p < 0.05$), indicating better risk-adjusted returns. The ability to disentangle risk and return as separate outcomes enables us to uncover this efficiency-improving behavior that is hidden by the aforementioned proxies, which the literature only reads as showing movements towards higher or lower risk & return positions together. Westerly or northwesterly movement on Figure A.14 would be entirely invisible to the standard proxies.

The treatment effects on portfolio risk are statistically robust but relatively modest in magnitude: portfolio standard deviation falls by 1.3% relative to the control mean. The first thing to say about this is that the treatment itself was not targeted to move these moments and that the amount of the transfer that reached agricultural production was commensurately small. The grant of approximately \$153 in 2017 USD represents only 4.4% of mean baseline household profit in the control group (\$3,482) and was not designed to alter crop portfolio composition. The portion of the grant that flowed into agricultural expenditure was smaller still: treated households increased input spending by an estimated \$31 (Table B.13) against a control-group mean of \$551, a 5.6% expansion of the input bundle. These portfolio-risk effects are the downstream consequence of that small input shift on crop selection and area allocation, not the response to a large agricultural shock.

Allocating an intercropped plot's area and input costs equally across its crops, rather than by revenue share, and rebuilding every outcome under that convention leaves the sign of every headline effect unchanged; the portfolio-SD effect is larger under the equal split, in dollars and relative to its control mean (Appendix C.1, Table C.1). Estimating $\mathbf{C}$ from community-year rather than farmer-year deviations gives a portfolio-SD effect of −\$3.62 per hectare ($p = 0.020$), against −\$4.01 from farmer-year deviations under the same specification, a gap of 0.21 standard errors (Appendix D.3; both under the headline specification).[20]

The unchanged per-hectare expected return in the MPT frame may at first appear inconsistent with the \$74 increase in *total* profit reported above, but the two findings are consistent. Total profit reflects hectares

[20] Reallocating cultivated *area* is the sharper test, since the crop weights w_c themselves depend on it. Appendix D.3 sets out the full menu of area-allocation alternatives (even-split area, sole-cropped-only $\mathbf{C}$, community-year aggregation) and why each carries its own bias. Appendix G separately tests sensitivity to the dip-test crop classification of Section 2.2, comparing treatment effects on high-intensity, low-intensity and merged crop shares (Table C.2).

cultivated × per-hectare expected return × realized state of the world; the grant operates primarily on the extensive margin (cultivated area increases by 0.36 HA) and on per-hectare *variance*, not on per-hectare expected return. The MPT framework is constructed in per-hectare terms so that scale-invariant portfolio moments can be compared across heterogeneous farm sizes. Higher total profit alongside unchanged per-hectare expected return is exactly the pattern predicted by an extensive-margin-driven portfolio expansion: farmers cultivate more land at similar per-hectare returns and slightly lower per-hectare variance, yielding more total income with less per-acre risk. Nor does the reconciliation hide an intensification margin: expenditure per hectare is flat (\$0.59, $p = 0.69$), and where input intensity does move it shows up as share shifts between the high- and low-input variants of maize and groundnut (Appendix C.2), which the portfolio measure already prices as separate assets.

6.2.2 Efficiency and the Channel of Risk Reduction

Table B.14 examines portfolio efficiency more directly. Households move \$1.48 closer to their crop-set specific efficient frontier ($p < 0.01$, control mean \$49), indicating improved portfolio efficiency. By contrast, the distance to the household's *cardinality-restricted* frontier (the closest point on the k-EF computed with k equal to the household's own number of cultivated crops) changes by only −\$0.2/HA ($p = 0.58$, control mean \$39). The efficiency gain at the unconstrained frontier therefore comes mainly from households growing more crops rather than from allocating better among the crops they already grew, a decomposition I return to below that answers the second of the two questions Section 5 closed on. It shows that households can close the gap by expanding portfolio size, rather than by reallocating land across the same number of crops.

The Sharpe ratio *at* the household's ideal point, meaning the closest portfolio on the efficient frontier, shows no statistically significant change (coefficient 0.09, $p = 0.39$, control mean 119.6). This is the level of risk-adjusted return available at an unconstrained point rather than its observed Sharpe ratio, so an unchanged value says the targeted point did not move. Read together, these two measures imply that grant recipients move *closer* to the efficient frontier without moving to a *different point along* it: the grant relaxes the constraints holding households inside the frontier without altering the risk-preferences-driven location they target on it. This is the signature pattern of constraint relaxation operating without any change in underlying preferences for risk. Preferences that shift with wealth cannot be fully excluded of course, but this \$153 cash transfer was unlikely to change wealth in a way that would have passed through to preferences.

The final two columns of Table B.14 isolate the channel through which the risk reduction operates: they are the Selection Ratio and Diversification Ratio measures introduced in Section 2.6 The self-insurance channel is quiet on average: the Selection Ratio shows no significant treatment effect (Table B.14), so grant recipients did not on average systematically shift land toward inherently lower-variance crops relative to

the most profitable crop available in their community. The reduction instead runs through the hedging channel. The Diversification Ratio rises by 1.16 points ($p < 0.01$), from a control mean of 135.4, indicating that grant recipients are better able to exploit imperfect correlations across crop returns to achieve greater risk-reduction benefits. The decomposition reveals the production cost of incomplete markets in this setting: without the grant, liquidity-constrained households accept higher portfolio variance, and the consumption volatility that comes alongside it because they cannot afford the crop combinations that provide better risk management. The grant partially releases this constraint, allowing households to achieve the same expected returns with less risk.

Structural models of incomplete markets predict that the inefficiency exists: Donovan (2021) and Rodríguez-Sala (2026) show that liquidity constraints push households away from efficiency, but they typically frame the inefficiency as households being forced *toward safer, lower-return* crops than they would otherwise choose, so that constraint relaxation should let them *take more risk* for higher returns. The empirical pattern here runs in the opposite direction: relaxing the liquidity constraint moves households toward *less* risk at the same level of expected return. The mechanism is not return-chasing but diversification: the grant enables entry into the right additional crops and the resulting wider crop set lowers portfolio variance through the off-diagonal terms of the variance-covariance matrix. The structural prediction of inside-the-frontier positioning is confirmed; the predicted direction of correction is not.

6.2.3 Crop-Level Margins and the Fixed-Cost Predictions

The portfolio-level treatment effects documented above are an aggregate consequence of shifts in the crops households choose to cultivate and how they allocate land across them. Input intensity within a crop is not one of those margins: expenditure per hectare did not move (\$0.59, $p = 0.69$), and where intensity does shift it is carried by the high- and low-input crop variants (Appendix C.2); Appendix C.8 confirms that the control-estimated moments price treated portfolios in both return and risk. Figure A.25 decomposes these shifts along three margins: whether a crop is grown, what share of land is allocated to it and how many hectares it occupies in levels. The left panel reports OLS estimates of the grant's effect on the extensive margin, the probability of cultivating each crop type. Grant recipients are significantly more likely to cultivate beans, rice, sorghum, the residual "other" category, high-intensity groundnut and okra, with Bambara groundnut significant at the ten percent level. That set spans the variance distribution rather than concentrating in the safer crops, which is the first sign that the risk reduction works through reallocation rather than through entry into low-variance crops. No crop type is less likely to be cultivated by treated farmers, and low-intensity maize is the only crop with a negative point estimate.

The middle panel reports fractional multinomial logit (FMLOGIT) marginal effects on the intensive mar-

gin, the share of cultivated land allocated to each crop. The FMLOGIT is required to interpret reallocation across crops coherently because shares are a set of zero-sum outcomes. Grant recipients allocate a larger share of land to high-intensity groundnut and to beans, moving land away from low-intensity varieties. The right panel then displays movement in levels. Grant recipients expand hectares across nearly every crop: 14 of the 16 point estimates are positive. The contrast between these two margins is informative. Hectares rise significantly for beans, high-intensity groundnut and okra, the broad-based expansion the added liquidity enables, yet the expansion does not translate into proportional share changes. Shares fall for the two low-intensity staples: low-intensity maize, the highest-share crop at baseline, though only suggestively at $p = 0.14$, and low-intensity groundnut ($p = 0.04$). They rise for beans and high-intensity groundnut, with sweet potato smaller and weaker; the coefficients are in Figure A.25. Households are expanding cultivation across the board, but the compositional rebalancing that drives the portfolio risk reduction points in a consistent direction: the two largest declines in share are both low-intensity varieties, and the two largest increases are beans and high-intensity groundnut. Despite expanding farm size in levels, farmers achieve a more efficient portfolio position; the broad expansion does not dilute efficiency but rather provides the additional acreage through which the compositional shift operates.

If the fixed-cost interpretation is correct, exogenous liquidity provision should operate primarily on the extensive margin. The grant will enable households to clear $L_{c,\min}$ for crops they previously did not cultivate or cultivated at the floor, rather than along the intensive margin of reweighting existing crops. Two empirical predictions follow. First, hectare treatment effects should concentrate on crops with small baseline shares, where new entry near $L_{c,\min}$ is the binding margin. Second, the extensive-margin response should be larger for households with enough baseline farmland to clear the floor with the grant in hand and smaller for households so liquidity-constrained that even a \$153 transfer leaves them unable to add new crops.

Figure A.26 reports the grant treatment effect on cultivated hectares for each crop, pooled within bins of the crop's baseline share in the control group. Expressed as a share of each bin's control-group mean, the hectare effect declines monotonically with baseline share, as the first prediction requires. It is 16.0% of the control-group mean for crops with baseline shares below 2%, 7.5% in the 2–5% bin, 4.2% in the 5–10% bin, and 3.3% for crops above 10%. The lowest bin holds 7 crops (Bambara, beans, fonio, okra, sesame, sweet potato, and yam). The effect is statistically significant at $p < 0.001$ for the lowest-baseline-share bin and weaker but consistently positive for higher bins. Figure A.27 tests the second prediction, reporting the same treatment effect restricted to the seven small-baseline-share crops, estimated separately within each baseline-farmland quintile. The effect on the smallest-farm quintile (Q1, median 2.15 HA) is statistically indistinguishable from zero (0.006 HA per crop, $p = 0.26$). Quintiles 2 through 5 show positive treatment effects between 0.010 and 0.023 HA per crop, significant at $p < 0.05$ in three of the four, corresponding to

11–21% of the bin's control-group mean.

This is the pattern Section 5.6 predicted and Section 5.5 left for the grant to settle: entry costs bind where land per crop is scarce, so cash adds crops on farms with room for them and cannot on farms without. The Q1 null is therefore not evidence against the fixed-cost story but a sharper version of it: the threshold binds these households so tightly that the $31 of incremental input spending the grant induces (Section 6) is insufficient to clear the cost of adding a marginal crop.

A separate concern is that the same Q1 null could also be generated by alternative mechanisms (soil heterogeneity that increases with farm size, labor supply, or learning), each of which predicts weaker responses on small farms for reasons unrelated to indivisibility. Appendix G.6 weighs the null against these alternatives and finds that neither fits.[21] This places a partial lower bound on the implicit threshold magnitude: the per-crop entry cost in this setting exceeds what $31 of additional working capital can finance for a 2-hectare farm. That points toward instruments calibrated to the per-crop entry threshold, such as input-package programs or working-capital credit sized to the floor, which could deliver gains where unconditional cash falls short for the smallest farms. Together, three pieces of evidence (the size-diversification gradient, a constrained efficient frontier whose gap is concentrated at the low-σ end, and grant treatment effects that concentrate on small-baseline-share crops and middle-quintile farms) support the cardinality-constrained interpretation of smallholder portfolio choice.

6.2.4 Distributional Effects

If, as Section 5.5 suggests, credit constraints mediate distance to the frontier, then the grant's effect should be systematically larger for the households furthest from it. The average effects mask meaningful variation across the outcome distribution. Quantile regressions (Figure A.28) show that the grant's effect on portfolio SD grows from approximately −$3.5 per hectare at the 10th percentile of the σ distribution to about −$7.8 per hectare at the 90th percentile, with the largest effects concentrated above the 70th percentile. The same pattern appears for distance to the efficient frontier and for the Sharpe ratio: high-distance and low-Sharpe households absorb the largest gains, consistent with a catch-up dynamic in which the grant disproportionately benefits the households operating furthest from their unconstrained optimum.

Taken together, these results suggest that liquidity constraints prevent households from achieving efficient diversification. When constraints are relaxed, households expand cultivation, add crops to their portfolios, and shift toward lower-variance crop combinations that better exploit the covariance structure of returns. The resulting portfolios achieve lower risk while maintaining expected returns, moving households closer to

[21] The size-diversification gradient does not vary with community-level baseline productivity dispersion (interaction $p = 0.500$), which rules out soil heterogeneity as the operative driver. It also flattens sharply below 1.50 HA of pre-floor cultivated area (slope-difference $p = 0.000$), a discontinuity the smooth alternatives would not produce.

their efficient frontiers and improving risk-adjusted welfare. The next subsection asks who delivers these effects, splitting them by baseline wealth.

6.3 Heterogeneity by Baseline Wealth

The headline ATE pools across the baseline wealth distribution, and the pooled estimate hides a pattern no single prediction anticipates. Estimating the headline ANCOVA within terciles of the baseline wealth index (Section 4.4; Table B.15), only the tails reduce portfolio risk: the poorest tercile's standard deviation falls by \$6.0 per hectare ($p < 0.10$) with a Sharpe gain of 1.0, the richest tercile's by \$5.1 ($p = 0.10$) with a Sharpe gain of 0.9, the middle tercile's not at all, and no tercile's expected return moves. The two textbook predictions each fit one tail. Under decreasing absolute risk aversion (Pratt, 1964; Rosenzweig and Binswanger, 1993) the response to a fixed transfer should fall smoothly with wealth; under a cash-on-hand threshold for indivisible investment (Fafchamps and Pender, 1997) it should kink, moving only the households the grant carries across. The tails reach the same destination by different routes, and the routes are what the rest of the table shows.

The poorest tercile reduces risk without expanding its crop-set. T1 households post the largest movements toward both the unconstrained efficient frontier (−\$2.3/HA, $p < 0.05$) and the cardinality-restricted frontier (−\$1.8/HA, $p < 0.05$), yet add no crops on net (+0.1, insignificant). The mechanism is reallocation rather than entry: land shares shift toward the household's lower-variance holdings without the portfolio growing in size. Treatment-induced crop switches within the poorest tercile raised the share of high-intensity groundnut while cutting the land share allocated to cotton, one of the highest-risk crops in the setting: the composition change names the crops through which risk fell. The Selection Ratio, the self-insurance margin, posts the largest movement of the three terciles (+2.8, $p = 0.11$) while the Diversification Ratio stays flat (+0.7, insignificant), the ordering a safer-crop channel implies.

The richest tercile instead reduces risk by expanding its crop set. T3 households add +0.4 unique crops per grant ($p < 0.01$) and raise their Sharpe ratio (+0.9, $p < 0.05$), and their portfolio-SD reduction (−\$5.1/HA) trends in the same direction while falling just short of conventional significance ($p = 0.10$). Their within-cardinality distances barely move (−\$0.3/HA, insignificant): for these households the risk reduction comes from the crop set widening rather than from optimizing inside a fixed one, and the crops they bring in behave as hedges against the rest of the portfolio. The channel decomposition again matches this explanation: the richest tercile's Diversification Ratio rises significantly (+1.2, $p < 0.05$) while its Selection Ratio does not, so within this group the added crops work as hedges that exploit the covariance structure rather than as substitutions toward inherently safer crops.

The middle tercile may hold a mix of households employing T1 or T3 style risk-reduction strategies that dilute the tercile average. Alternatively, the middle may be spending the risk budget that hedging frees up. It adds covariance-friendly crops while shifting land toward higher-variance holdings, so portfolio risk nets to zero. It shows the same extensive-margin response as T3 (+0.4 crops, $p < 0.01$) and a significant rise in the Diversification Ratio (+1.4, $p < 0.05$), yet converts neither into a reduction in portfolio risk, distance to the frontier, or improvement in the Sharpe ratio. Either way, the middle emerges as the tercile where the grant expands and rebalances the portfolio without changing its risk. Across the three terciles the ratio decompositions point the same way as the composition evidence without reaching conventional significance, and the between-tercile contrasts are themselves imprecise (Table B.15, right panel): the experiment was stratified on village rather than wealth, and only the contrast in portfolio breadth is separately distinguishable. The channel pattern is the suggested shape the estimates take, not a significant contrast.

Altogether, this pattern inverts the classic risk-bearing logic. The reversal points to the binding margin being a per-crop fixed entry cost rather than risk-bearing capacity. CRRA preferences and an indivisible per-crop fixed cost F can rationalize all three terciles. At the top, where households have enough cultivated area to absorb the indivisibility, the grant clears the fixed cost and funds additional hedging crops. At the bottom the grant is too small, so households reshuffle land instead of adding crops. The within-cardinality reallocation at the bottom and the extensive-margin entry at the top are two envelopes of the same (γ, F) structure, and which one dominates depends on where in the wealth distribution the fixed cost binds. Shenoy (2018) found evidence for the intensive margin alone in Thai rice; the Mali grant identifies both margins simultaneously, on different outcomes, in a setting where the per-crop fixed cost is binding.

The direction of this response is the mirror image of what index insurance produces, and the contrast sharpens the paper's central measurement point. Insurance interventions relax households' capacity to bear risk ex-ante, and the documented response is to take on more of it: Karlan et al. (2014)'s Ghanaian farmers narrow their portfolios to specialize in high-intensity maize, and insured farmers in Mobarak and Rosenzweig (2013) and Cai (2016) similarly shift toward what the authors describe as higher-return, higher-risk production technologies. The grant studied here moves households in the opposite direction, toward lower portfolio variance, because the constraint it relaxes is different: insurance lowers the price of bearing risk, while the grant absorbs the fixed costs of risk management strategies. Whether a transfer raises or lowers a smallholder's portfolio risk is therefore not a property of the transfer but of the friction it removes. The contrast also shows why a count of crops cannot stand in for a measure of risk: insurance recipients who specialize and grant recipients who diversify move in opposite directions in risk-return space, and only a covariance-based measure recognizes that both are optimizing against the constraint they face.

Appendix F interacts the grant with continuous baseline moderators and finds all 20 linear interactions

null (Table F.1), which is what a both-tails pattern implies rather than evidence against it.

7 Economic Magnitude of the Risk Reduction

To put the average risk reduction in context, moving along the efficient frontier to achieve an equivalent 1.3% reduction in portfolio SD would require foregoing approximately $4.07 per hectare in expected return. Grant recipients achieve this risk reduction while sacrificing only $2.38 per hectare (statistically insignificant), consistent with the fall in distance to the frontier (Table B.14) rather than with movement along it.[22] As an additional dollar benchmark for the magnitude of portfolio inefficiency, the median pure-control household at baseline sits with an approximately $66 per hectare gap to the empirical efficient frontier evaluated at its own portfolio standard deviation, equivalent to about $412 per year at that sample's mean cultivated area of 6.3 hectares. This is the vertical analogue of the dollar-return-loss statistic used in Calvet et al. (2007).[23] The cash transfer's $4.01-per-hectare reduction in portfolio standard deviation again reflects movement of households toward the efficient frontier rather than along it. The unrealized headroom of roughly $66 per hectare is the ceiling for what a treatment that fully alleviates all binding constraints, bringing each household to the frontier, could achieve.

Translating the variance reduction into a dollar welfare gain requires a coefficient of relative risk aversion, γ, which the lottery survey questions from this experiment estimate too noisily to use directly. I work around this by adapting the frontier-projection methodology of Calvet et al. (2007): for each baseline farmer, I project their observed (σ_h, μ_h) position onto their crop-set specific efficient frontier in two directions, West (same μ, lower σ) and North (same σ, higher μ), shown by Figure A.29. The tangency condition from the mean-variance utility function optimization, $\gamma = (\partial\mu/\partial\sigma)/\sigma$, at each projection point pins down a farmer-level implied γ under the maintained assumption that the farmer was optimizing at the projected frontier point. The West projection identifies the most-risk-averse interpretation of the farmer's position (γ_{west}), the North projection the most-risk-tolerant (γ_{north}); together they bound the farmer's true γ under mean-variance preferences. Reading a single risk-aversion parameter off a farmer's observed portfolio position would not be appropriate here: the evidence throughout this paper is that observed positions are constrained, reflecting displacement by binding constraints at least as much as preferences. Applying the headline ATEs to each farmer's γ bound and averaging gives a CE welfare gain of 5.2% of grant value at the risk-tolerant lower

[22] The expected-return estimate carries a standard error of 1.76, so its ninety-five percent interval runs from −5.82 to 1.07 per hectare and contains the $4.07 along-frontier price. On that estimate alone, movement purely along the frontier cannot be rejected. Evaluated at the two ends of the interval, the certainty-equivalent bound of Appendix G.10 spans -11 to -4% of grant value at the lower end and 21 to 28% at the upper. What separates the two readings is the second margin rather than precision on $\Delta\mu$: a household giving up more than the along-frontier price would be moving away from the frontier, and its measured distance to the frontier falls instead.

[23] I report the vertical mean-gap rather than their Sharpe-loss version because no defensible risk-free rate exists in this setting (Section 2.1).

bound and 12% at the risk-averse upper bound.[24]

This range is a calibration rather than an estimate: its two endpoints come from two risk-aversion bounds applied to the same treatment-effect point estimates, so its width reflects uncertainty about preferences and not statistical precision.[25] What the range does establish is that the sign of the gain does not depend on where within it preferences sit: the certainty-equivalent gain remains positive for any risk aversion above $\gamma^* = 0.00194$, which lies a factor of 1.5 below the risk-tolerant bound and 2.2 below the risk-averse one.

The projection methodology has a limitation in that under binding constraints, the household's observed (σ, μ) position reflects both unobserved risk preferences and constraint severity. This means that the projection identifies a composite of the two rather than γ alone. Households with more inefficient positions will have wider ranges of potential γ values than households with relatively more efficient positions. It is natural to ask whether the implied preferences are plausible given the projection reads each household's risk aversion from its position. The frontier positions imply a CRRA coefficient of roughly 0.9 to 1.4 after being translated to the scale of relative risk aversion. This range is inside the conventional development-literature range of one to five. The projection recovers preferences of conventional magnitudes without having assumed them, a test it could have failed. In spite of this, the identification limitation above still binds as the implied coefficients remain influenced by both preferences and constraint severity. The projection reads part of the constraint severity as preference. Appendix G.10 develops this diagnostic and reports the smaller welfare gain that conventional CRRA values would imply.

8 Conclusion

Three main findings emerge from applying Modern Portfolio Theory to this unconditional cash grant RCT of Malian smallholders. The first is that exogenous cash grants relax an implied constraint through the extensive margin on average. Among crops with small baseline shares cultivation expands by 16.0% of the control mean while high-share crops change little, and the smallest-farm quintile fails to clear the threshold even with the grant. Next, the resulting reallocation from the grant reduced portfolio standard deviation by 4.01 USD per hectare (1.3% of the control mean) but left expected profit per hectare unchanged. This combination improves the Sharpe ratio and moves households closer to their crop-set specific efficient frontiers, yielding

[24] Full derivation and computation in Appendix G.10. Imposing the maintained hypothesis $\Delta\mu = 0$, which the ANCOVA cannot reject ($p = 0.177$), raises the bound to 16–23% of grant value; the point-estimate version reported here is therefore the conservative one. The $\Delta\mu$ point estimate is negative but statistically indistinguishable from zero ($p = 0.177$), and remains insignificant under a crop-set fixed effect with 18 levels that absorbs cross-community heterogeneity in crop availability ($p = 0.112$; see Appendix G.8). Using the point estimate rather than imposing $\Delta\mu = 0$ is the conservative choice, since the negative estimate trims the welfare gain at both bounds.

[25] The calculation propagates the point estimates of $\Delta\mu$ and $\Delta\sigma^2$ without carrying their sampling uncertainty, and the figures should accordingly be read as a magnitude calibration rather than as a tested claim. The underlying risk reduction is itself statistically significant, but the calibration adds a translation into welfare units, and that translation introduces additional parameter uncertainty that the treatment-effect estimates do not have.

certainty-equivalent welfare gains of 5.2–12% of grant value. Finally, the heterogeneity by baseline wealth showcases multiple mechanisms working across the distribution. The poorest households reallocate land within their existing portfolio size and still move toward the frontier while the richest households expand their portfolios to reduce risk exposure. The bigger-picture contribution of these findings is twofold. As measurement, the MPT framework supplies portfolio-level outcomes that capture the risk consequences of crop choice in a way counts and concentration indices cannot. As explanation, its structural ingredients support a cardinality-constrained extension that accounts for the size-diversification gradient, the asymmetric extensive-margin effects, and the non-monotone wealth gradient.

These results have several implications for development policy. Cash transfers improve welfare not only through direct consumption effects but also by enabling better risk management, a channel largely absent from the existing literature on unconditional transfers. Because consumption smoothing is far from complete in rural West Africa (Townsend, 1994; Kazianga and Udry, 2006), reductions in agricultural income risk of this size plausibly carry through to consumption volatility as well. Cash transfers partially relieve the costly self-insurance households practice through their crop choices. But the gap that remains is large: at its current portfolio standard deviation, the median pure-control household forgoes roughly $66 per hectare of return relative to the efficient frontier (Section 7). Formal insurance, or instruments aimed directly at the entry cost, could therefore deliver gains that unconditional liquidity does not (Mobarak and Rosenzweig, 2013; Cai, 2016). Conversely, policies that encourage specialization should be weighed against the diversification benefits households mechanically sacrifice when they reduce the scope of their portfolios. For risk-averse smallholders without access to formal insurance, maintaining a diversified crop portfolio is a first-order risk-management strategy that specialization policies disrupt. Promoting monocropping without ensuring the target crop provides a superior risk-to-return ratio would be a textbook example of the "development mismatch" akin to what McGuirk and Nunn (2024) document when agricultural projects were placed in traditionally pastoral areas.

Several limitations of this analysis warrant discussion before concluding. First, the panel I use only spans four agricultural seasons (2010–2012, 2017). This means that rare but severe shocks may be underrepresented in my estimated return distributions relative to how they are weighted in households' mental models of weather and price shocks. Mali's four waves cannot settle this from the inside: with a single window there is no longer-horizon benchmark to compare against and no alternative windows to vary. Appendix D therefore turns to the seven-wave Uganda LSMS-ISA panel, where estimating the covariance matrix from all $\binom{7}{4} = 35$ four-wave subsets shows how much a four-year estimate is sensitive to which four years it happens to draw. I repeat the exercise within each of Uganda's 4 administrative regions, a further 140 covariance matrices, because the Mali sample sits in a single agroecological zone and pooling across zones would manufacture

instability that does not apply to it. The exercise shows that dependence is concentrated rather than general: among pairs where both crops are widely cultivated, four-wave estimates are centered on the seven-wave value with a root mean squared deviation of 0.15, against 0.33 where either crop is rare. Those widely cultivated crops carry the majority of portfolio weights in Mali, so I feel comfortable that a four-year window pins down the part of the covariance matrix the analysis rests on and leaves the sampling noise where it carries little weight.

Second, I do not observe risk preferences directly and must rely on an estimated measure for this dimension of farmers' responses. My ideal-point proxy for risk preferences is constructed from the efficient frontier. As the efficient frontier is itself a function of my estimated variance-covariance matrix the proxy and the portfolio risk measure it is compared against are not independent of one another (Appendix D.2 and Appendix D.1 stress-test that matrix against shrinkage and alternative estimation windows). The lottery elicitations available in these data are too noisy to recover γ (Section 5.5), so a better-designed preference instrument fielded alongside the experiment would be the natural remedy in future work. The endowment correlates of Section 5.5 provide suggestive evidence that constraints bind, but separating preferences from constraints more definitively would require additional data.

Finally, the MPT framework assumes households optimize over the first two moments of the return distribution: mean and variance and not skewness or kurtosis. If households exhibit prudence or temperance, behaviors requiring higher order moments, my frontier analysis is incomplete. Appendix G.4 examines the higher moments of crop returns directly, and I leave fuller preference extensions for future work.

The broadest point this paper makes is that counting unique crop varieties was never a sufficient measure of risk; it was a stand-in, adopted because estimating the variance-covariance matrix of crop returns requires plot-crop-level panel data that most studies did not have reason to collect. With that object estimated though, the crop portfolio stops being a metaphor: it is an observable risk-management technology smallholders already own, and development interventions can now be evaluated for what they do in tandem with it. As development interventions increasingly target agricultural productivity, understanding their full welfare effects requires accounting for changes in risk exposure, effects that standard outcome measures systematically miss.

References

Nurudeen Abdul Rahman, Asamoah Larbi, Bekele Kotu, Mary Otiwaa Asante, Daniel Brain Akakpo, Shaibu Mellon-Bedi, and Irmgard Hoeschle-Zeledon. Maize–legume strip cropping effect on productivity, income, and income risk of farmers in northern ghana. *Agronomy Journal*, 113(2):1574–1585, 2021. doi: https://doi.org/10.1002/agj2.20536. URL https://acsess.onlinelibrary.wiley.com/doi/abs/10.1002/agj2.20536.

Treb Allen and David Atkin. Volatility and the gains from trade. *Econometrica*, 90(5):2053–2092, 2022. doi: https://doi.org/10.3982/ECTA14411. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA14411.

Oriana Bandiera, Robin Burgess, Narayan Das, Selim Gulesci, Imran Rasul, and Munshi Sulaiman. Labor markets and poverty in village economies. *The Quarterly Journal of Economics*, 132(2):811–870, 2017.

Ziv Bar-Shira, Richard E. Just, and David Zilberman. Estimation of farmers' risk attitude: An econometric approach. *Agricultural Economics*, 17(2–3):211–222, 1997.

Lori Beaman, Dean Karlan, Bram Thuysbaert, and Christopher Udry. Selection into credit markets: Evidence from agriculture in mali. *Econometrica*, 91(5):1595–1627, 2023. doi: 10.3982/ECTA18916.

Dimitris Bertsimas and Romy Shioda. Algorithm for cardinality-constrained quadratic optimization. *Computational Optimization and Applications*, 43(1):1–22, 2009. doi: 10.1007/s10589-007-9126-9.

Michael J. Best and Robert R. Grauer. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: Some analytical and computational results. *The Review of Financial Studies*, 4(2):315–342, 1991. doi: 10.1093/rfs/4.2.315.

Steven C. Blank. The robustness of single index models in crop markets: A multiple index model test. *Western Journal of Agricultural Economics*, 16(2):259–267, 1991.

Oscar Burbano-Figueroa, Alexandra Sierra-Monroy, Adriana David-Hinestroza, Cory Whitney, Christian Borgemeister, and Eike Luedeling. Farm-planning under risk: An application of decision analysis and portfolio theory for the assessment of crop diversification strategies in horticultural systems. *Agricultural Systems*, 199:103409, 2022.

Jing Cai. The impact of insurance provision on household production and financial decisions. *American Economic Journal: Economic Policy*, 8(2):44–88, 2016. ISSN 19457731, 1945774X. URL http://www.jstor.org/stable/24739217.

Laurent E. Calvet, John Y. Campbell, and Paolo Sodini. Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy*, 115(5), 2007. doi: 10.1086/524204. URL https://doi.org/10.1086/524204.

John Y. Campbell. Household finance. *Journal of Finance*, 61(4):1553–1604, 2006. doi: 10.1111/j.1540-6261.2006.00883.x.

Calogero Carletto, Sydney Gourlay, and Paul Winters. From guesstimates to GPStimates: Land area measurement and implications for agricultural analysis. *Journal of African Economies*, 24(5):593–628, 2015. doi: 10.1093/jae/ejv011.

T.-J. Chang, N. Meade, J. E. Beasley, and Y. M. Sharaiha. Heuristics for cardinality constrained portfolio optimisation. *Computers & Operations Research*, 27(13):1271–1302, 2000. doi: 10.1016/S0305-0548(99)00074-X.

Raj Chetty. A new method of estimating risk aversion. *American Economic Review*, 96(5):1821–1834, 2006. doi: 10.1257/aer.96.5.1821.

Yves Choueifaty and Yves Coignard. Toward maximum diversification. *Journal of Portfolio Management*, 35(1):40–51, 2008.

Robert A. Collins and Peter J. Barry. Risk analysis with single-index portfolio models: An application to farm planning. *American Journal of Agricultural Economics*, 68(1):152–161, 1986. doi: 10.2307/1241659.

Leandro De Magalhães and Raül Santaeulàlia-Llopis. The consumption, income, and wealth of the poorest: An empirical analysis of economic inequality in rural and urban sub-saharan africa for macroeconomists. *Journal of Development Economics*, 134:350–371, 2018. doi: 10.1016/j.jdeveco.2018.05.014.

Stefan Dercon. Risk, crop choice, and savings: Evidence from tanzania. *Economic Development and Cultural Change*, 44(3):485–513, 1996. doi: 10.1086/452234.

Salvatore Di Falco and Jean-Paul Chavas. On crop biodiversity, risk exposure, and food security in the highlands of ethiopia. *American Journal of Agricultural Economics*, 91(3):599–611, 2009.

Kevin Donovan. The equilibrium impact of agricultural risk on intermediate inputs and aggregate productivity. *The Review of Economic Studies*, 88(5):2275–2307, 2021. doi: 10.1093/restud/rdaa084.

Dennis Egger, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker. General equilibrium effects of cash transfers: Experimental evidence from Kenya. *Econometrica*, 90(6):2603–2643, 2022.

Nadja El Benni and Robert Finger. Where is the risk? price, yield and cost risk in Swiss crop production. *Revue d'Études en Agriculture et Environnement*, 95(3):299–326, 2014. doi: 10.4074/S1966960714013022.

Marcel Fafchamps and John Pender. Precautionary saving, credit constraints, and irreversible investment: Theory and evidence from semiarid india. *Journal of Business & Economic Statistics*, 15(2):180–194, 1997.

FEWS NET. Seasonal calendar for a typical year: Mali. Famine Early Warning Systems Network, USAID, June 2024. URL https://fews.net/west-africa/mali/seasonal-calendar/july-2024.

Andrew D. Foster and Mark R. Rosenzweig. 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, 2022. doi: 10.1086/717890.

Rudolf J. Freund. The introduction of risk into a programming model. *Econometrica*, 24(3):253–263, 1956. doi: 10.2307/1911630.

Douglas Gollin and Christopher Udry. Heterogeneity, measurement error, and misallocation: Evidence from african agriculture. *Journal of Political Economy*, 129(1):1–80, 2021. doi: 10.1086/711369. URL https://doi.org/10.1086/711369.

Arne Hallam. Economies of size and scale in agriculture: An interpretive review of empirical measurement. *Applied Economic Perspectives and Policy*, 13(1):155–172, 1991. doi: https://doi.org/10.2307/1349565. URL https://onlinelibrary.wiley.com/doi/abs/10.2307/1349565.

J. A. Hartigan and P. M. Hartigan. The dip test of unimodality. *The Annals of Statistics*, 13(1):70–84, 1985.

Johannes Haushofer and Jeremy Shapiro. The short-term impact of unconditional cash transfers to the poor: Experimental evidence from kenya. *The Quarterly Journal of Economics*, 131(4):1973–2042, 2016. doi: 10.1093/qje/qjw025.

Peter B. R. Hazell. A linear alternative to quadratic and semivariance programming for farm planning under uncertainty. *American Journal of Agricultural Economics*, 53(1):53–62, 1971. doi: 10.2307/1237838.

Peter B. R. Hazell and Roger D. Norton. *Mathematical Programming for Economic Analysis in Agriculture*. Macmillan, New York, 1986.

Earl O. Heady. Diversification in resource allocation and minimization of income variability. *Journal of Farm Economics*, 34(4):482–496, 1952. ISSN 10711031. URL http://www.jstor.org/stable/1233230.

John Heaton and Deborah Lucas. Portfolio choice and asset prices: The importance of entrepreneurial risk. *Journal of Finance*, 55(3):1163–1198, 2000. doi: 10.1111/0022-1082.00244.

N. J. Jobst, M. D. Horniman, C. A. Lucas, and G. Mitra. Computational aspects of alternative portfolio selection models in the presence of discrete asset choice constraints. *Quantitative Finance*, 1(5):489–501, 2001. doi: 10.1088/1469-7688/1/5/301.

Richard E. Just and Rulon D. Pope. Stochastic specification of production functions and economic implications. *Journal of Econometrics*, 7(1):67–86, 1978. ISSN 0304-4076. doi: https://doi.org/10.1016/0304-4076(78)90006-4. URL https://www.sciencedirect.com/science/article/pii/0304407678900064.

Dean Karlan, Robert Osei, Isaac Osei-Akoto, and Christopher Udry. Agricultural Decisions after Relaxing Credit and Risk Constraints *. *The Quarterly Journal of Economics*, 129(2):597–652, 02 2014. ISSN 0033-5533. doi: 10.1093/qje/qju002. URL https://doi.org/10.1093/qje/qju002.

Harounan Kazianga and Christopher Udry. Consumption smoothing? livestock, insurance and drought in rural burkina faso. *Journal of Development Economics*, 79(2):413–446, 2006.

Phoebe Koundouri, Marita Laukkanen, Sami Myyrä, and Céline Nauges. The effects of EU agricultural policy changes on farmers' risk attitudes. *European Review of Agricultural Economics*, 36(1):53–77, 2009.

Takashi Kurosaki and Marcel Fafchamps. Insurance market efficiency and crop choices in pakistan. *Journal of Development Economics*, 67(2):419–453, 2002. doi: 10.1016/S0304-3878(01)00198-2.

Rachid Laajaj and Karen Macours. The complexity of multidimensional learning in agriculture. *American Economic Review*, 2024. Forthcoming.

Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004. doi: 10.1016/S0047-259X(03)00096-4.

Remy Levin and Daniela Vidart. The yeoman's portfolio: Measuring historical risk preferences using crop choice. Working Paper 35634, National Bureau of Economic Research, 2026.

Haim Levy and Harry M. Markowitz. Approximating expected utility by a function of mean and variance. *American Economic Review*, 69(3):308–317, 1979.

Chunjie Li, Tjeerd-Jan Stomph, David Makowski, Haigang Li, Chaochun Zhang, Fusuo Zhang, and Wopke van der Werf. The productive performance of intercropping. *Proceedings of the National Academy of Sciences*, 120(2):e2201886120, 2023. doi: 10.1073/pnas.2201886120.

Elaine M. Liu. Time to change what to sow: Risk preferences and technology adoption decisions of cotton farmers in china. *The Review of Economics and Statistics*, 95(4):1386–1403, 2013. doi: 10.1162/REST_a_00295.

Clifton Makate, Arild Angelsen, Stein Terje Holden, and Ola Tveitereid Westengen. Crops in crises: Shocks shape smallholders' diversification in rural ethiopia. *World Development*, 159:106054, 2022. ISSN 0305-750X. doi: https://doi.org/10.1016/j.worlddev.2022.106054. URL https://www.sciencedirect.com/science/article/pii/S0305750X22002443.

Clifton Makate, Arild Angelsen, Stein Terje Holden, and Ola Tveitereid Westengen. Evolution of farm-level crop diversification and response to rainfall shocks in smallholder farming: Evidence from malawi and tanzania. *Ecological Economics*, 205:107736, 2023. ISSN 0921-8009. doi: https://doi.org/10.1016/j.ecolecon.2022.107736. URL https://www.sciencedirect.com/science/article/pii/S0921800922003974.

Harry Markowitz. Portfolio selection*. *The Journal of Finance*, 7(1):77–91, 1952. doi: https://doi.org/10.1111/j.1540-6261.1952.tb01525.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.1952.tb01525.x.

Eoin F. McGuirk and Nathan Nunn. Development mismatch? evidence from agricultural projects in pastoral africa. Working Paper 33191, National Bureau of Economic Research, 2024.

David McKenzie. Beyond baseline and follow-up: The case for more t in experiments. *Journal of Development Economics*, 99(2):210–221, 2012. ISSN 0304-3878. doi: https://doi.org/10.1016/j.jdeveco.2012.01.002. URL https://www.sciencedirect.com/science/article/pii/S030438781200003X.

Richard O. Michaud. The markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, 45(1):31–42, 1989. doi: 10.2469/faj.v45.n1.31.

Jeffrey D. Michler and Anna L. Josephson. To specialize or diversify: Agricultural diversity and poverty dynamics in ethiopia. *World Development*, 89:214–226, 2017.

Ahmed Mushfiq Mobarak and Mark R. Rosenzweig. Informal risk sharing, index insurance, and risk taking in developing countries. *American Economic Review*, 103(3):375–80, May 2013. doi: 10.1257/aer.103.3.375. URL https://www.aeaweb.org/articles?id=10.1257/aer.103.3.375.

Jonathan Morduch. Income smoothing and consumption smoothing. *Journal of Economic Perspectives*, 9 (3):103–114, September 1995. doi: 10.1257/jep.9.3.103. URL https://www.aeaweb.org/articles?id=10.1257/jep.9.3.103.

National Agricultural Statistics Service. Farms and land in farms, 2022 summary. Technical report, United States Department of Agriculture, 2023.

David M. G. Newbery and Joseph E. Stiglitz. *The Theory of Commodity Price Stabilization: A Study in the Economics of Risk*. Clarendon Press, Oxford, 1981.

Raphaël Paut, Rodolphe Sabatier, and Marc Tchamitchian. Reducing risk through crop diversification: An application of portfolio theory to diversified horticultural systems. *Agricultural Systems*, 168:123–130, 2019. ISSN 0308-521X. doi: https://doi.org/10.1016/j.agsy.2018.11.002. URL https://www.sciencedirect.com/science/article/pii/S0308521X18306802.

John W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1-2):122–136, 1964.

Md. Raseduzzaman and Erik Steen Jensen. Does intercropping enhance yield stability in arable crop production? a meta-analysis. *European Journal of Agronomy*, 91:25–33, 2017. ISSN 1161-0301. doi: https://doi.org/10.1016/j.eja.2017.09.009. URL https://www.sciencedirect.com/science/article/pii/S1161030117301399.

Albert Rodríguez-Sala. Incomplete markets, crop portfolios, and agricultural productivity. Working paper, 2026.

Mark R. Rosenzweig and Hans P. Binswanger. Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments. *The Economic Journal*, 103(416):56–78, 01 1993. ISSN 0013-0133. doi: 10.2307/2234337. URL https://doi.org/10.2307/2234337.

Laura Schechter. Risk aversion and expected-utility theory: A calibration exercise. *Journal of Risk and Uncertainty*, 35(1):67–76, 2007. doi: 10.1007/s11166-007-9017-6.

William F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk*. *The Journal of Finance*, 19(3):425–442, 1964. doi: https://doi.org/10.1111/j.1540-6261.1964.tb02865.x. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.1964.tb02865.x.

Ajay Shenoy. Risky income or lumpy investments? evidence on two theories of underspecialization. *Economic Development and Cultural Change*, 66(4):629–669, 2018. doi: 10.1086/696176.

Tavneet Suri. Selection and comparative advantage in technology adoption. *Econometrica*, 79(1):159–209, 2011. doi: 10.3982/ECTA7749.

Robert M. Townsend. Risk and insurance in village india. *Econometrica*, 62(3):539–591, 1994.

Peter H. Westfall and S. Stanley Young. *Resampling-Based Multiple Testing: Examples and Methods for p-Value Adjustment*. John Wiley & Sons, New York, 1993.

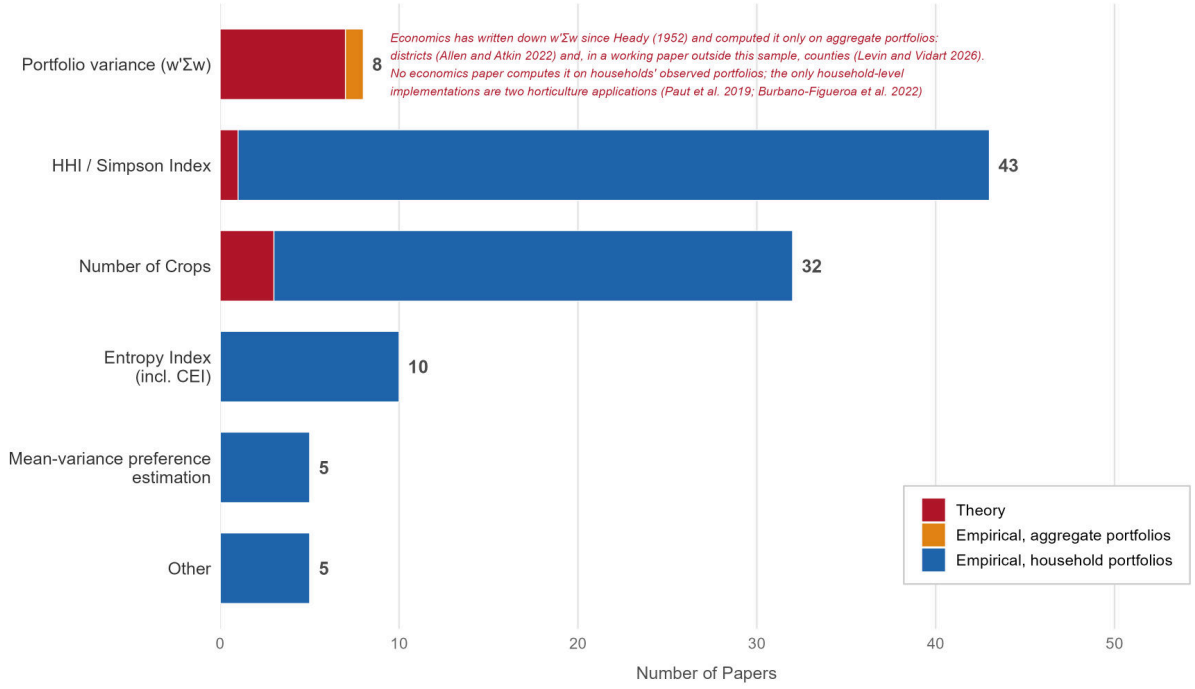
Jinpu Wu, Xingguo Bao, Jiudong Zhang, Binglin Lu, Weiping Zhang, Ragan M. Callaway, and Long Li. Temporal stability of productivity is associated with complementarity and competitive intensities in intercropping. *Ecological Applications*, 33(1):e2731, 2023. doi: https://doi.org/10.1002/eap.2731. URL https://esajournals.onlinelibrary.wiley.com/doi/abs/10.1002/eap.2731.

Frederic J. Zimmerman and Michael R. Carter. Asset smoothing, consumption smoothing and the reproduction of inequality under risk and subsistence constraints. *Journal of Development Economics*, 71(2): 233–260, 2003. doi: 10.1016/S0304-3878(03)00028-2.

Appendix A: Figures

Figures referenced in the main text.

Figure A.1: Measures of Crop Diversification in the Economics Literature



Notes: Systematic review of economics journal articles measuring crop diversification since Heady (1952); working papers, books, and non-economics outlets excluded (74 articles; classification rules and per-paper resolutions in the replication materials). Papers may use multiple measures. The portfolio-variance bar collects papers that derive $w'\Sigma w$ analytically and the one journal article that computes it on aggregate portfolios, at the district level (Allen and Atkin, 2022); Levin and Vidart (2026) do the same at the county level in a working paper outside the journal sample. None computes it on households' observed portfolios, so the bar carries no household-level empirical mass. The mean-variance preference-estimation bar collects papers that use MV machinery to recover risk attitudes from acreage choices.

Figure A.2: Heterogeneity in risk preferences along Capital Market Line

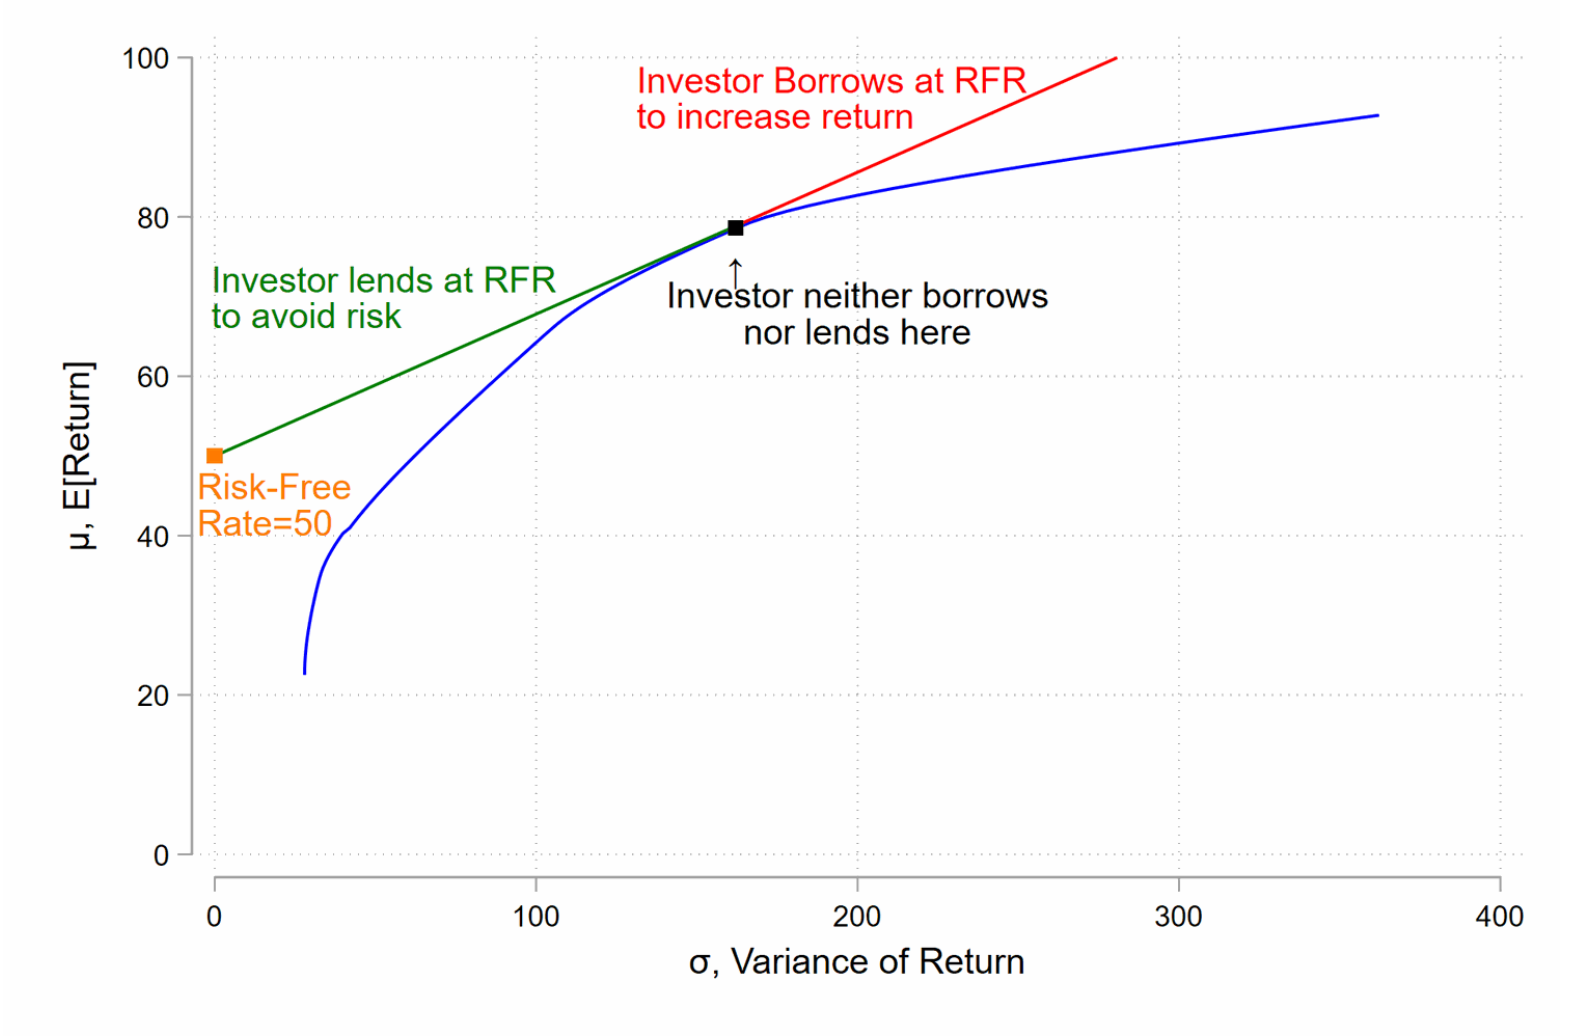


Figure A.3: Heterogeneity in risk preferences along efficient frontier

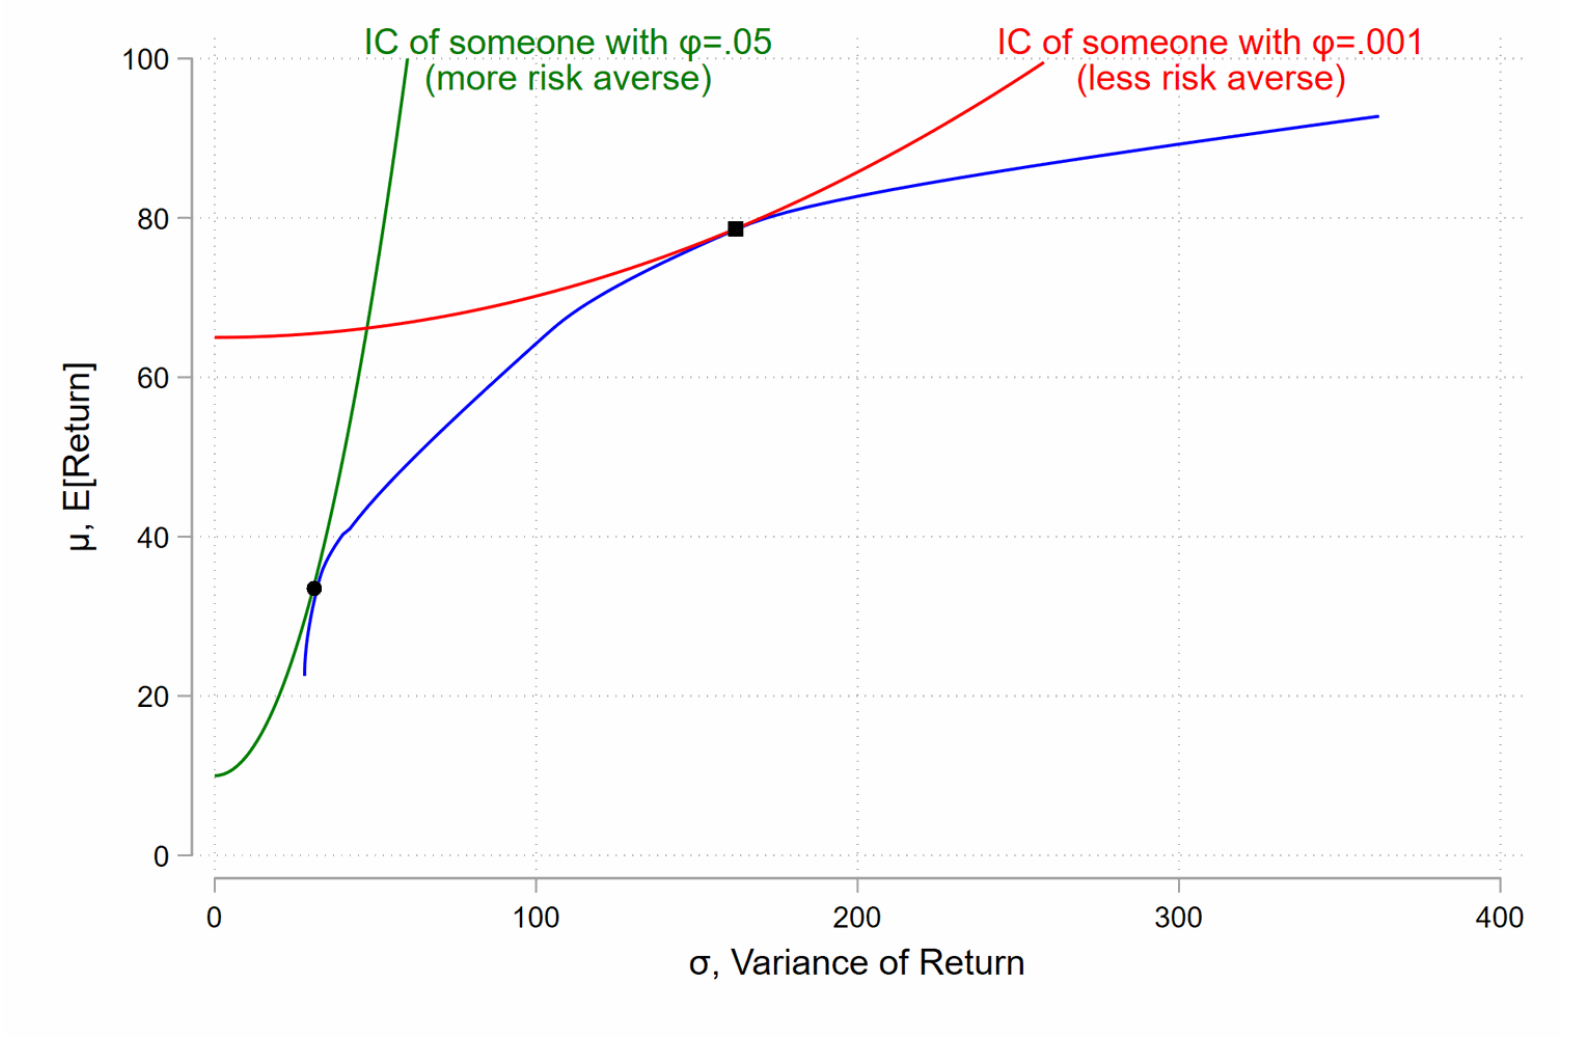
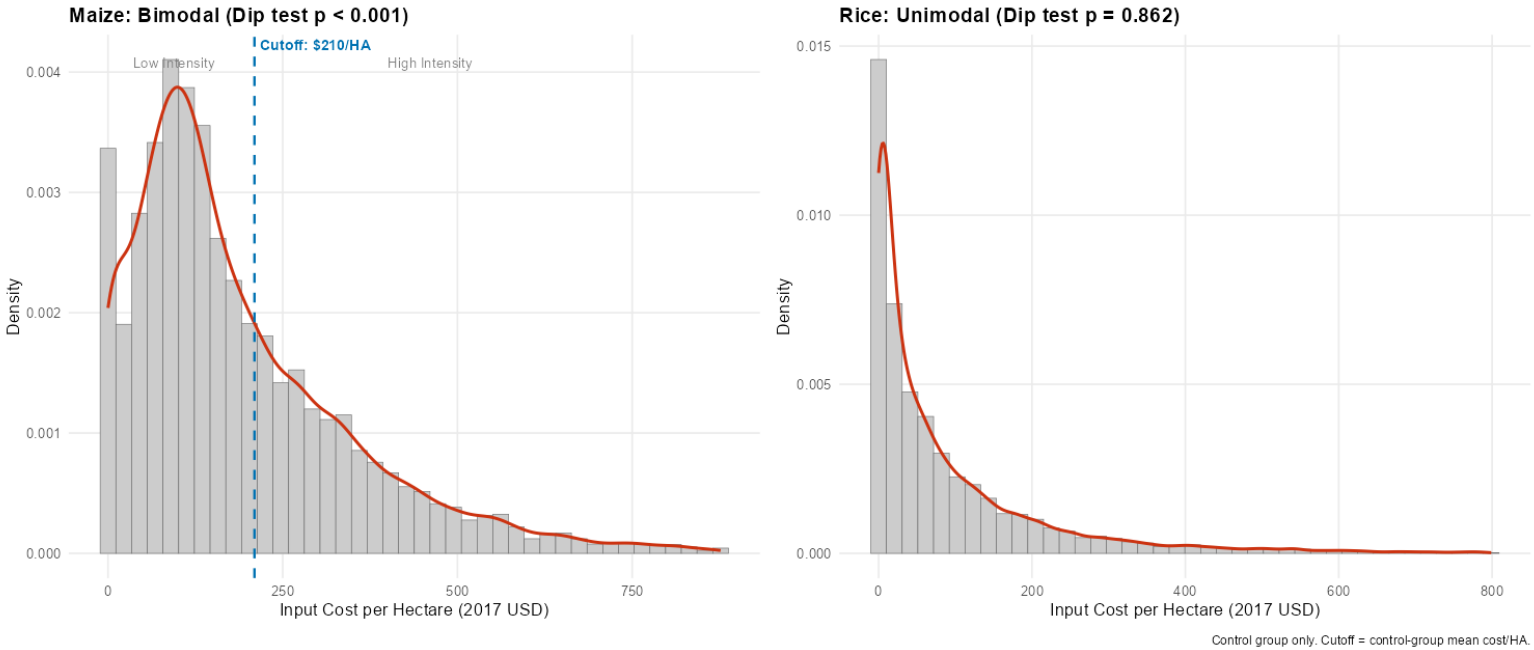
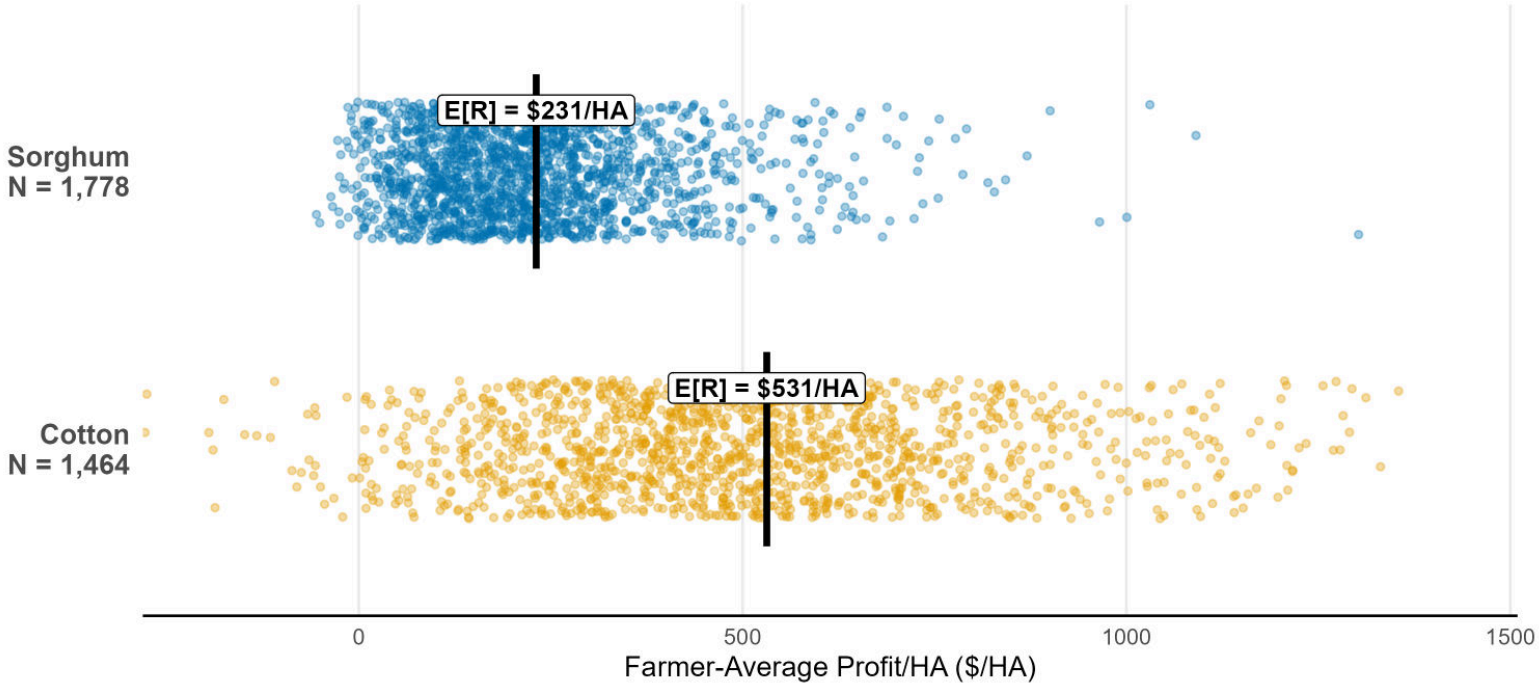


Figure A.4: Hartigan's Dip Test of Unimodality on Input Cost per Hectare



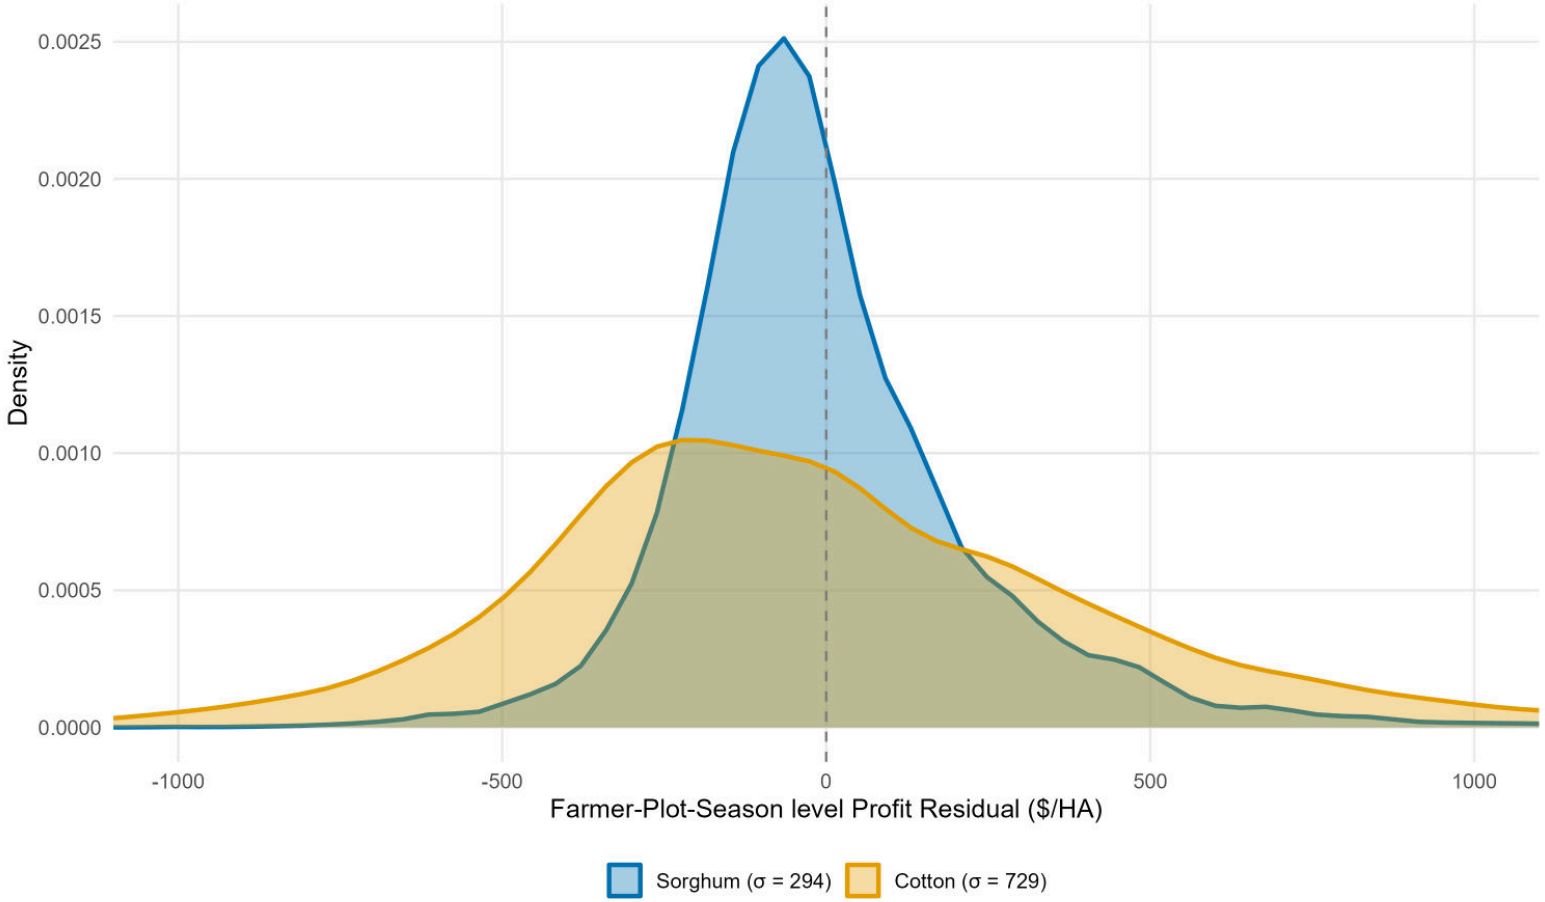
Notes: Distribution of input costs per hectare for maize (left) and rice (right). Hartigan's Dip Test rejects unimodality for maize ($p < 0.001$) and does not reject for rice ($p = 0.86$), so maize is split into high- and low-intensity variants and rice is treated as a single asset class. The test's null hypothesis is that a single mode describes the distribution; a rejection therefore indicates only that one mode does not, and does not establish how many regimes are present. The dashed line is the control-group mean cost per hectare, the cutoff at which the split is imposed; it is a chosen discretization of a continuous intensity margin rather than an estimated boundary between technologies. Control group observations only.

Figure A.5: Crop Expected Returns: Sorghum vs. Cotton



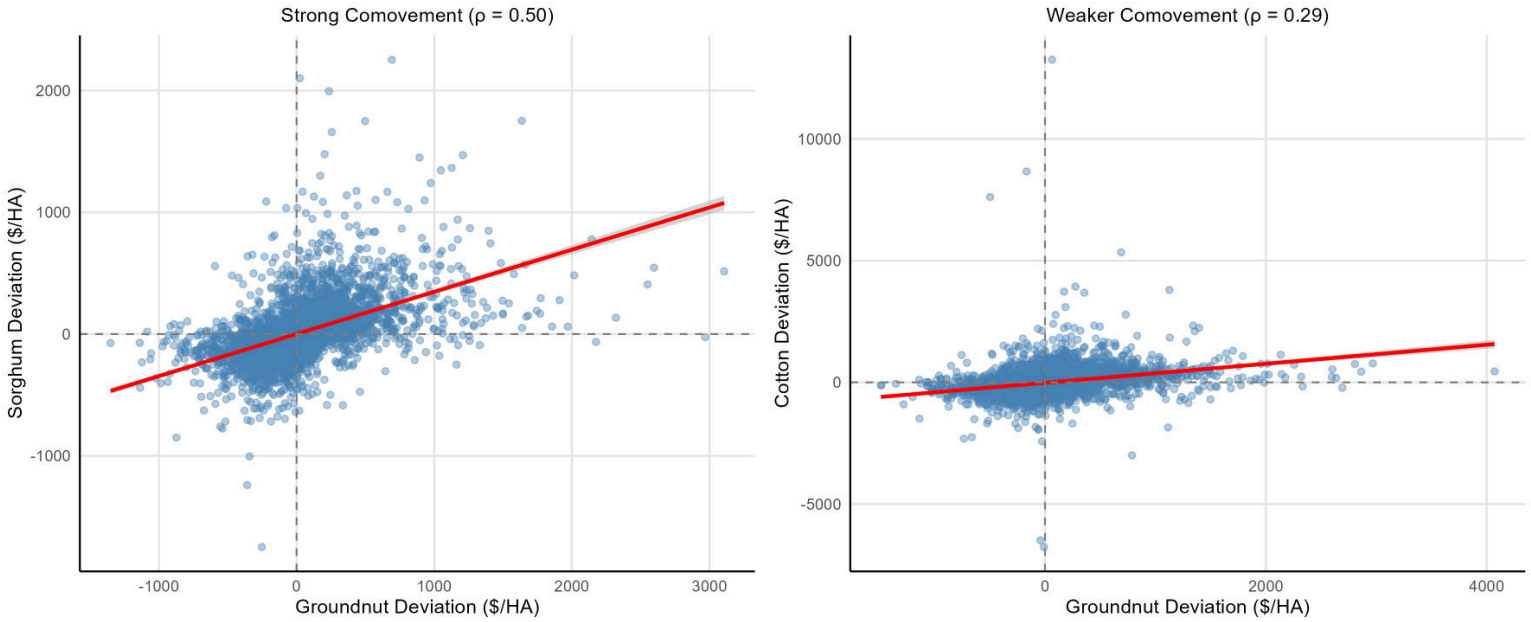
Notes: Each dot is one farmer's average profit per hectare across the years they cultivated the crop. Vertical bars mark the cross-farmer mean, which I use as the expected return $\bar{R}_c$ for each crop. Sorghum ($232/HA) is a representative low-return staple; cotton ($531/HA) is a representative high-return cash crop. The visible spread within each crop is across-farmer heterogeneity, not the within-farmer variance used to estimate σ_c.

Figure A.6: Crop Profit Variability: Sorghum vs. Cotton



Notes: Density of farmer-year profit residuals after removing each farmer's own crop-specific mean. The width of each density is the within-farmer, across-year variance from which σ_c is estimated. Cotton residuals have a much wider distribution ($\sigma = 729$/HA) than sorghum ($\sigma = 294$/HA), illustrating the cross-crop heterogeneity in idiosyncratic risk that the variance-covariance matrix $\mathbf{C}$ is built from.

Figure A.7: Crop Profit Comovement: Scatter of Profit Deviations Across Crop Pairs



Notes: Scatter plots of farmer-year profit deviations for two crop pairs. Left panel: groundnut (low) and sorghum exhibit strong positive comovement ($\rho = 0.50$ across farmer-year deviations), meaning cultivating both provides limited diversification benefit. Right panel: groundnut (low) and cotton exhibit weaker comovement ($\rho = 0.29$), offering greater scope for risk reduction through joint cultivation. Each point represents a farmer-year observation.

Figure A.8: First two moments of estimated crop returns

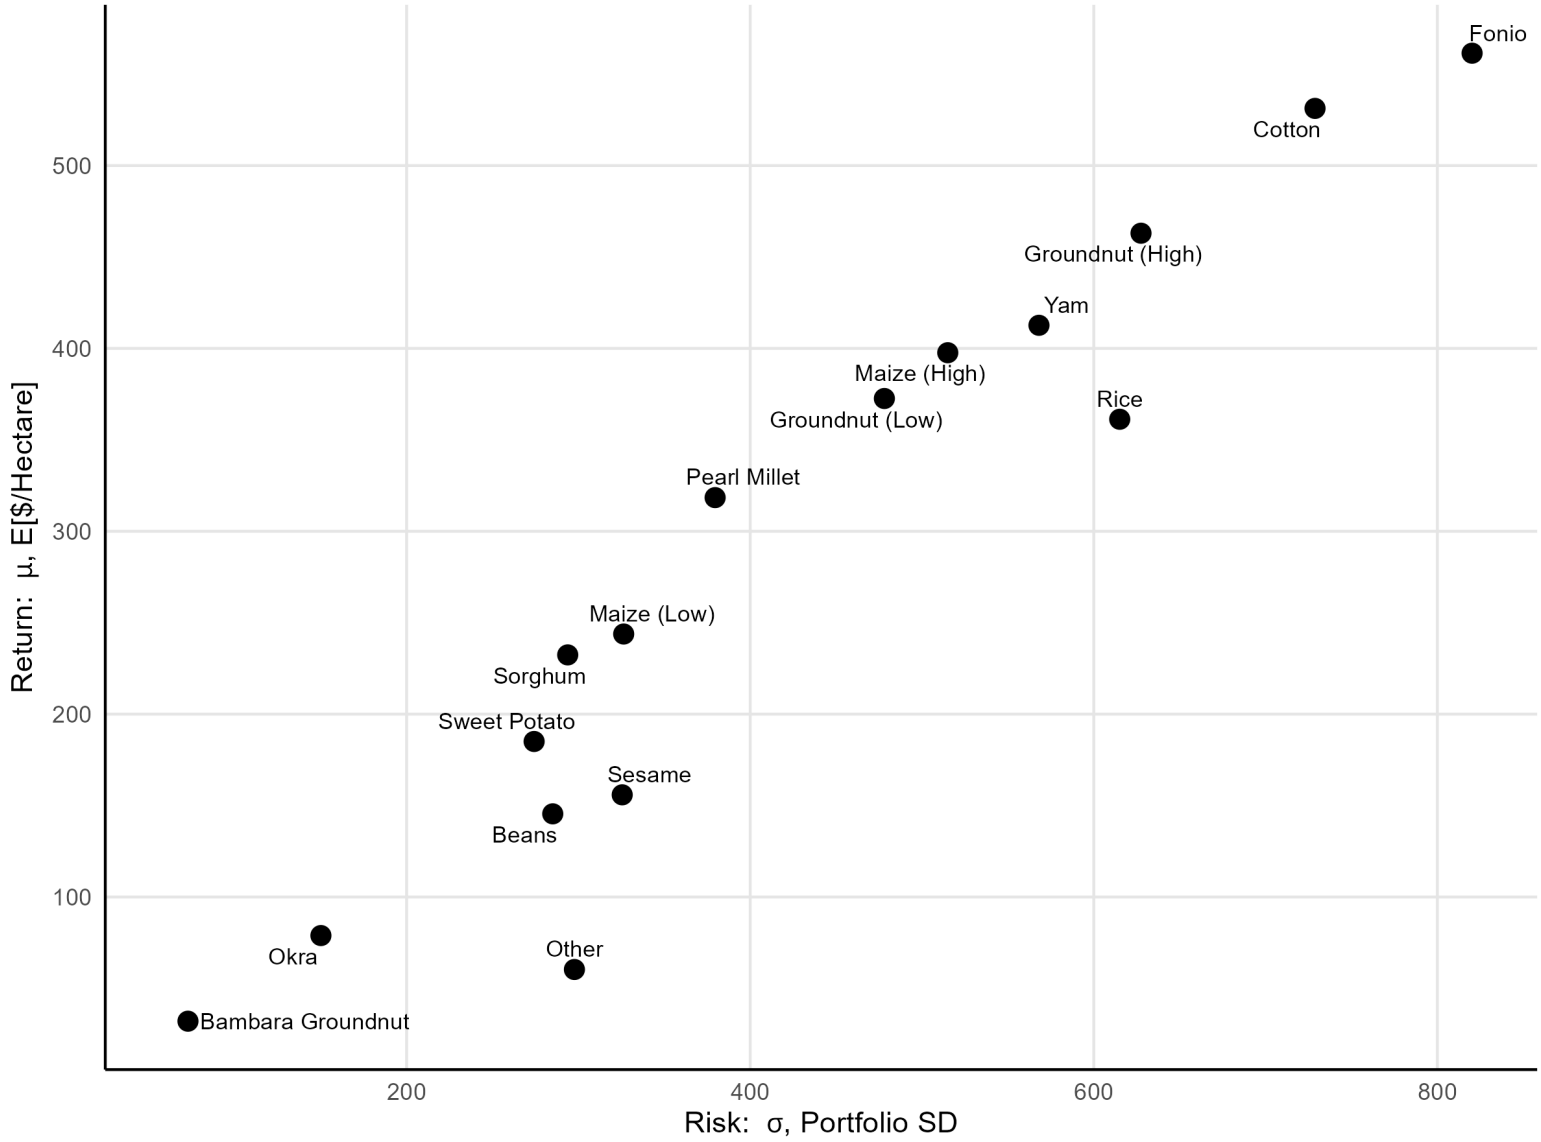


Figure A.9: Sharpe Ratio Diagram

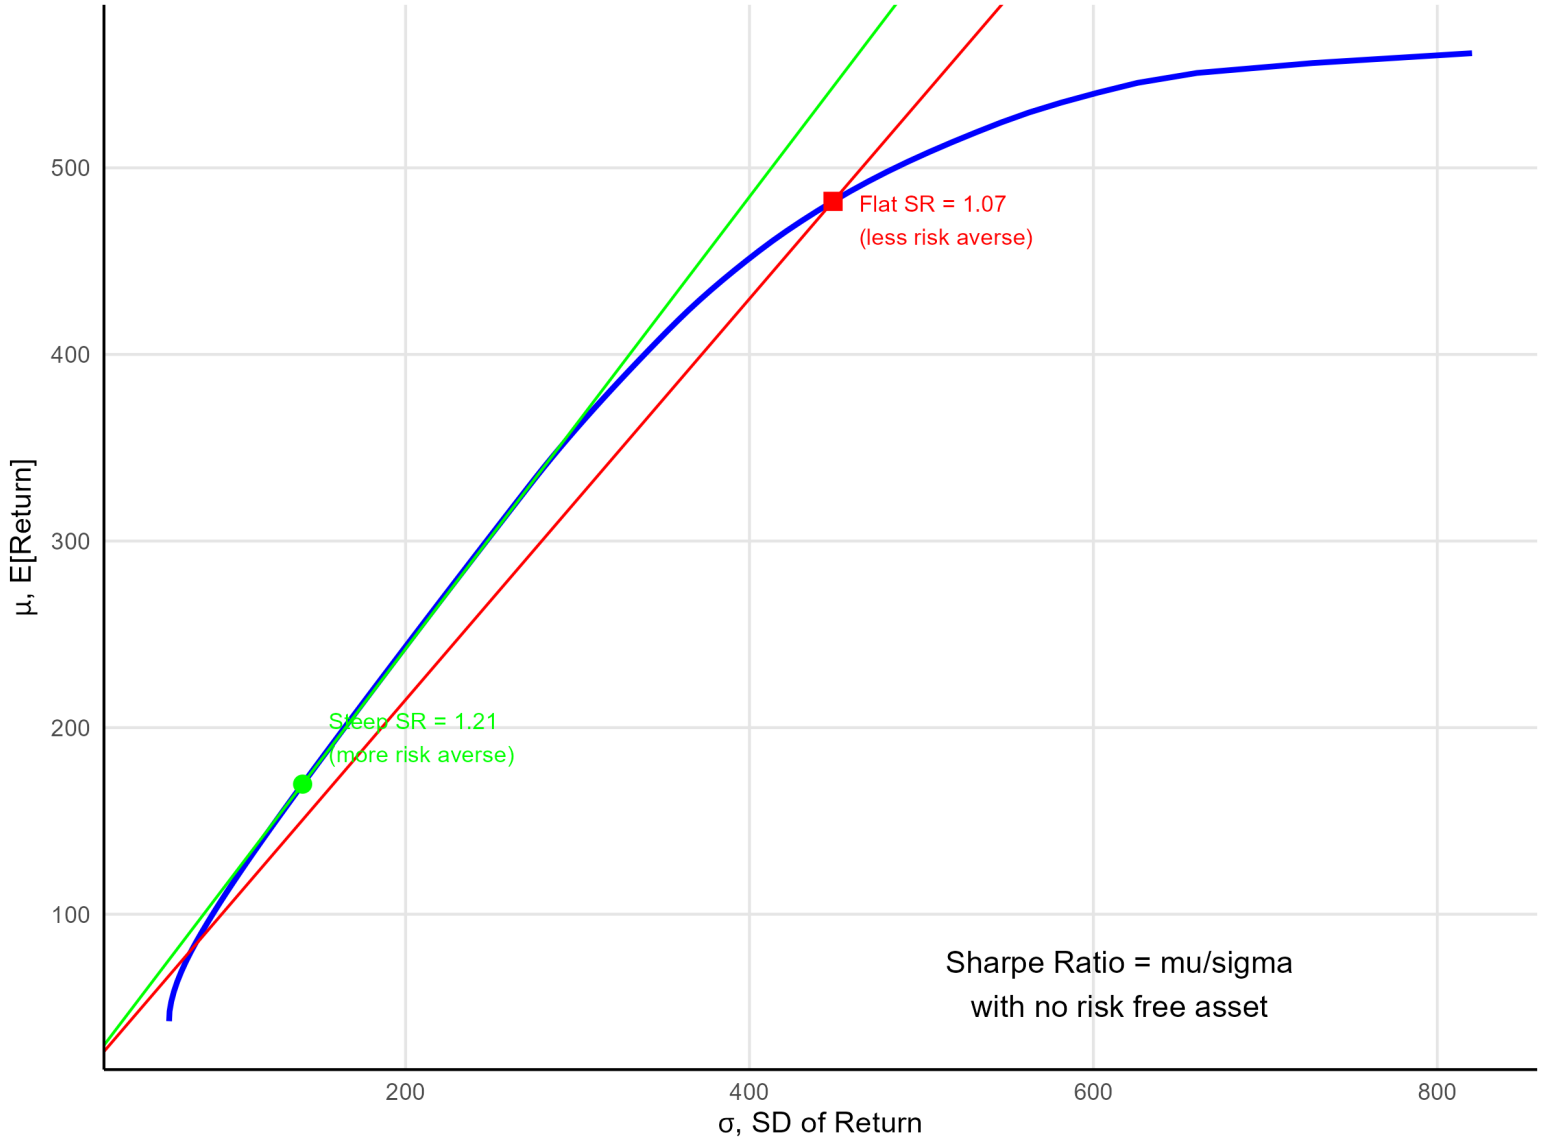
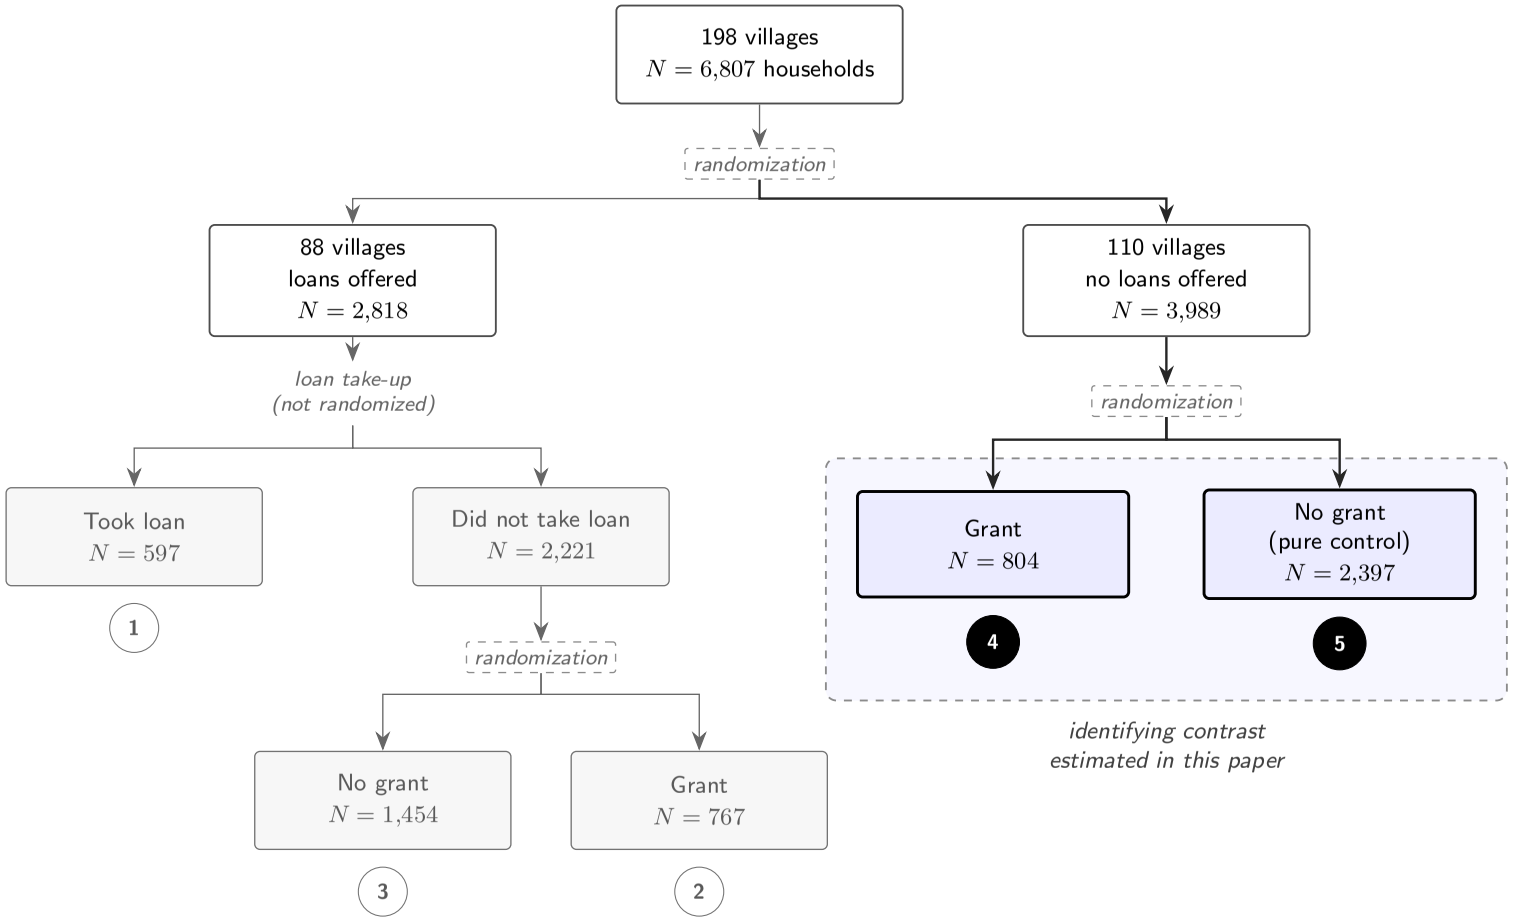
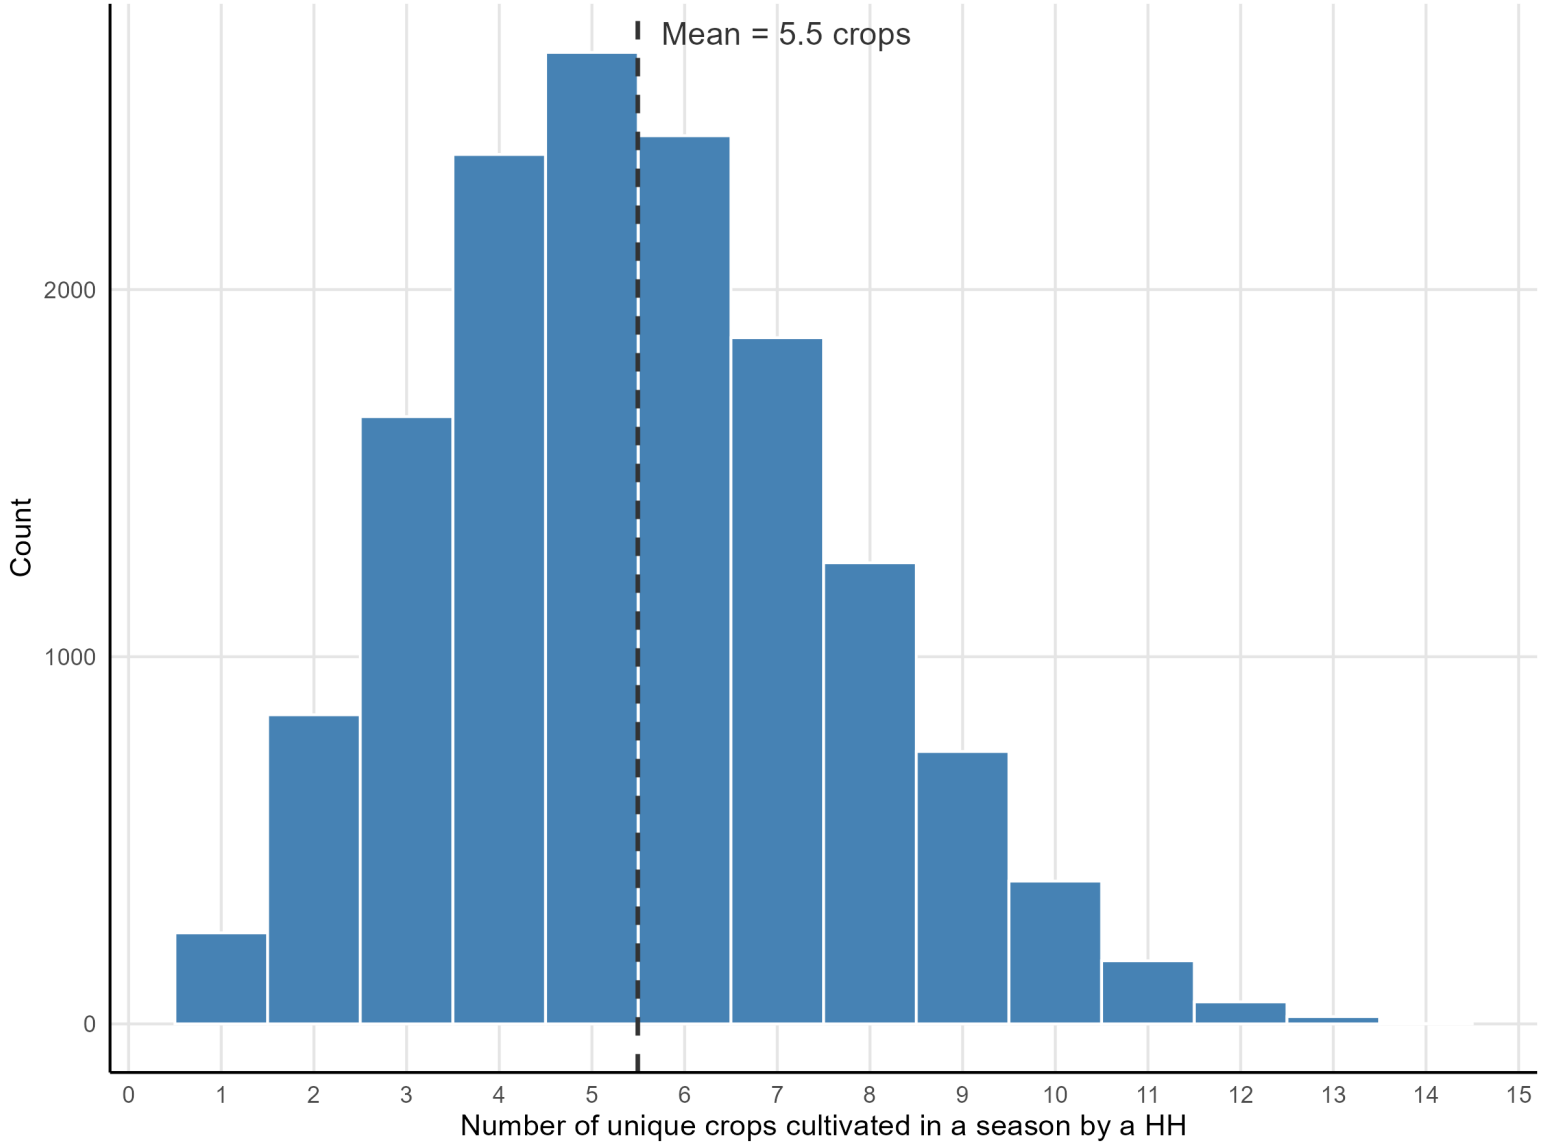


Figure A.10: Experimental Treatment Structure



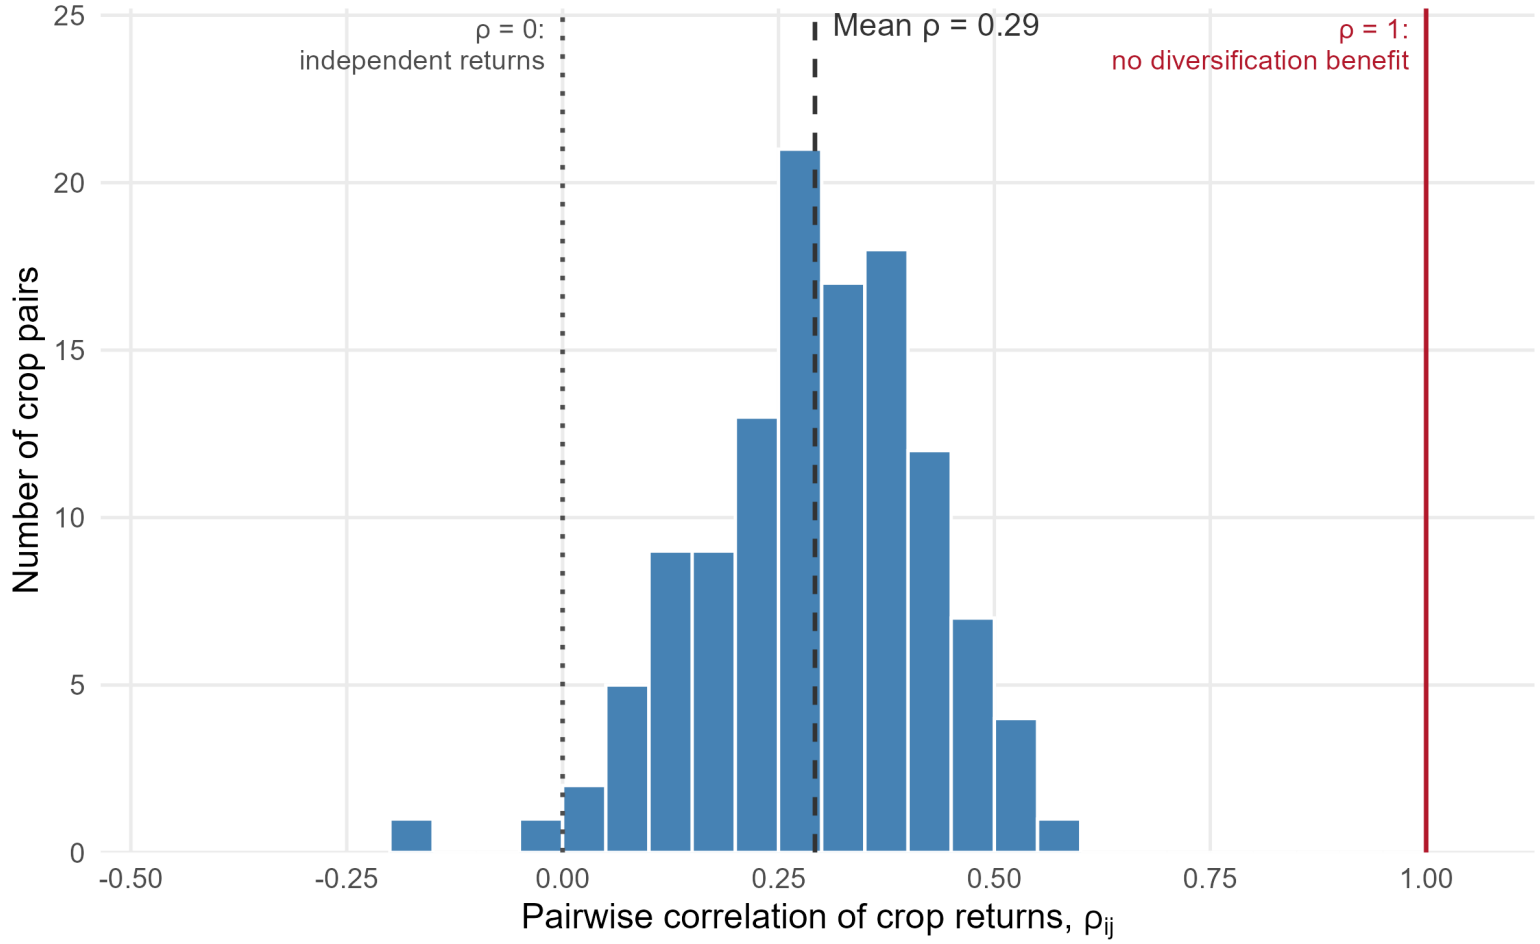
Notes: The original experiment cross-cut grants and loans across villages, producing the five arms numbered here. This paper uses arms 4 (grant villages, no loans offered) and 5 (pure control) to isolate the effect of an unconditional cash transfer, shaded above. Sample restricted to balanced panel households present in all survey rounds. Counts are the experimental design totals reported by Beaman et al. (2023); this paper's estimation sample is described in Section 4.2.

Figure A.11: Number of Crops Planted in A Farming Season



Notes: An observation is one household-year: 14,686 household-years across the four survey rounds. The sample is the 110 no-loan villages, pooling grant recipients (arm 4) and pure control (arm 5) across all four rounds; households from the loan villages and the other experimental arms are excluded. The dashed line marks the mean of 5.5 crops. Monocropping is rare in this setting: 248 household-years, 1.7% of the panel, cultivate a single crop.

Figure A.12: Distribution of Pairwise Crop-Return Correlations



Notes: Each observation is one of the 120 unique crop pairs in the estimated correlation matrix, computed from farmer-year profit deviations. The dashed line marks the mean of 0.29. A pair at $\rho = 1$ would offer no diversification benefit; every pair below it offers some. Only 2 of the 120 pairs are negative, so diversification value in this setting comes from imperfect comovement rather than from true hedges.

Figure A.13: Crop Asset Properties and Efficient Frontier

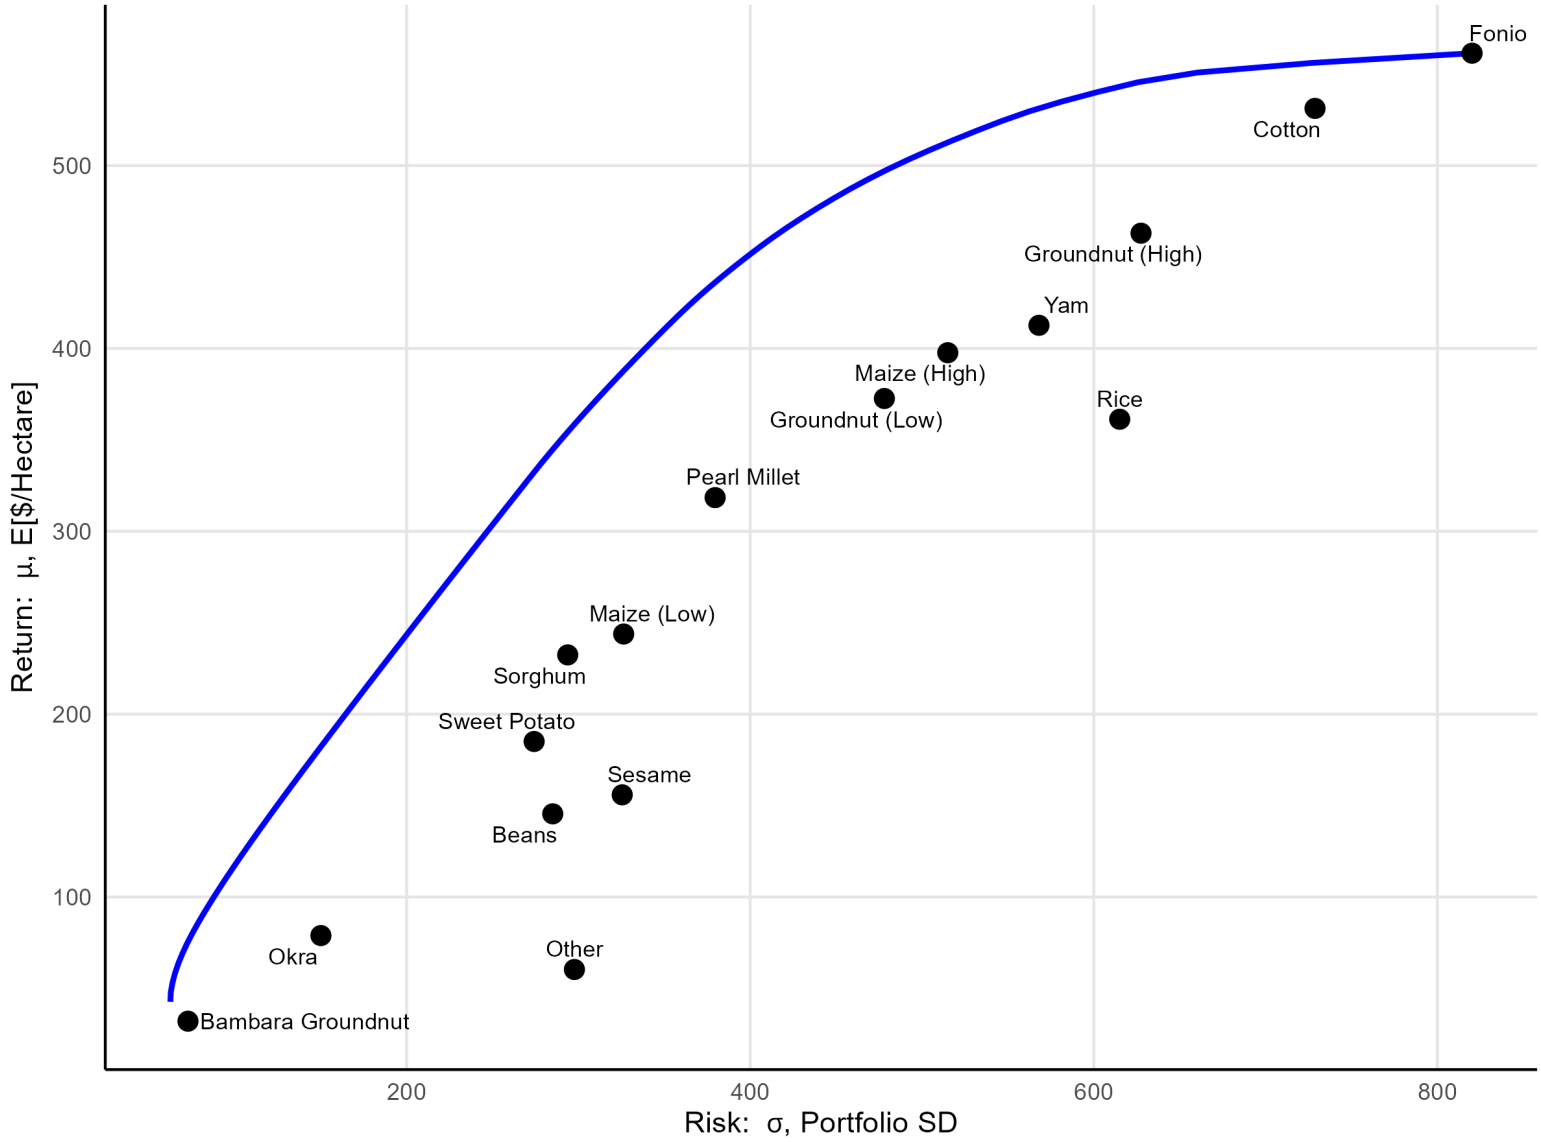
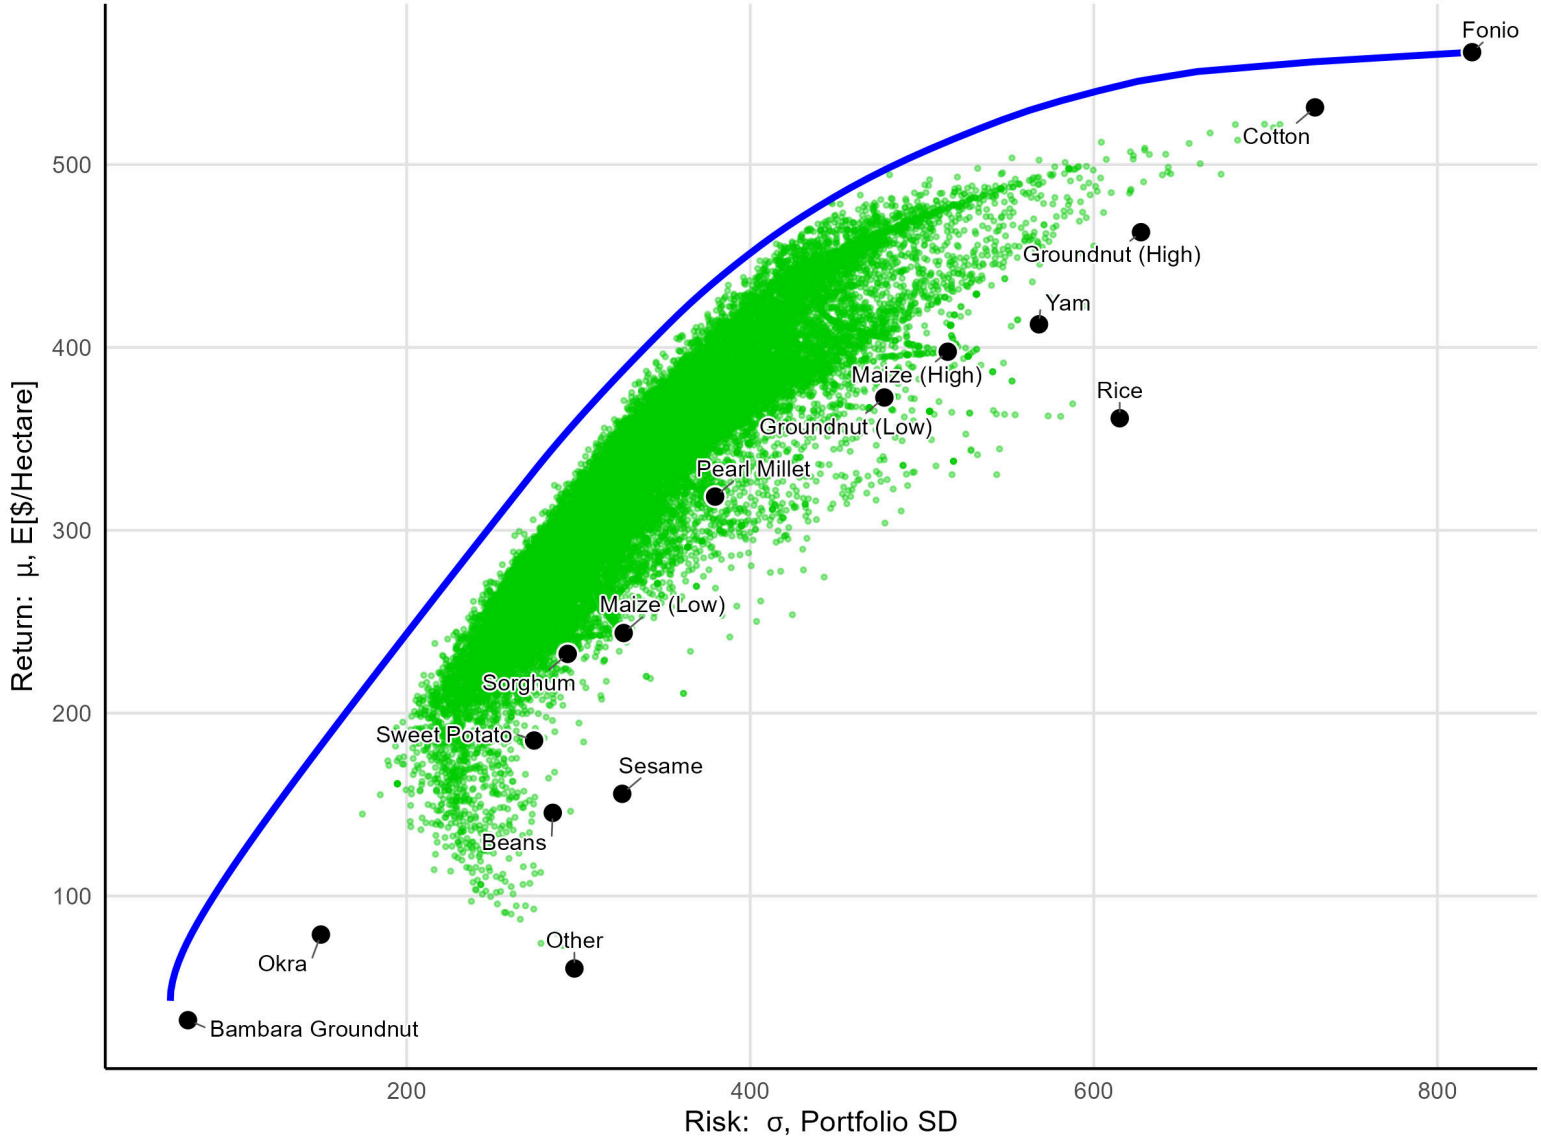


Figure A.14: Household Portfolio Properties



Notes: Green dots are observed household-year portfolios; the blue curve is the efficient frontier for the full menu of 16 crop types.

Figure A.15: Crop-Set Specific Efficient Frontiers

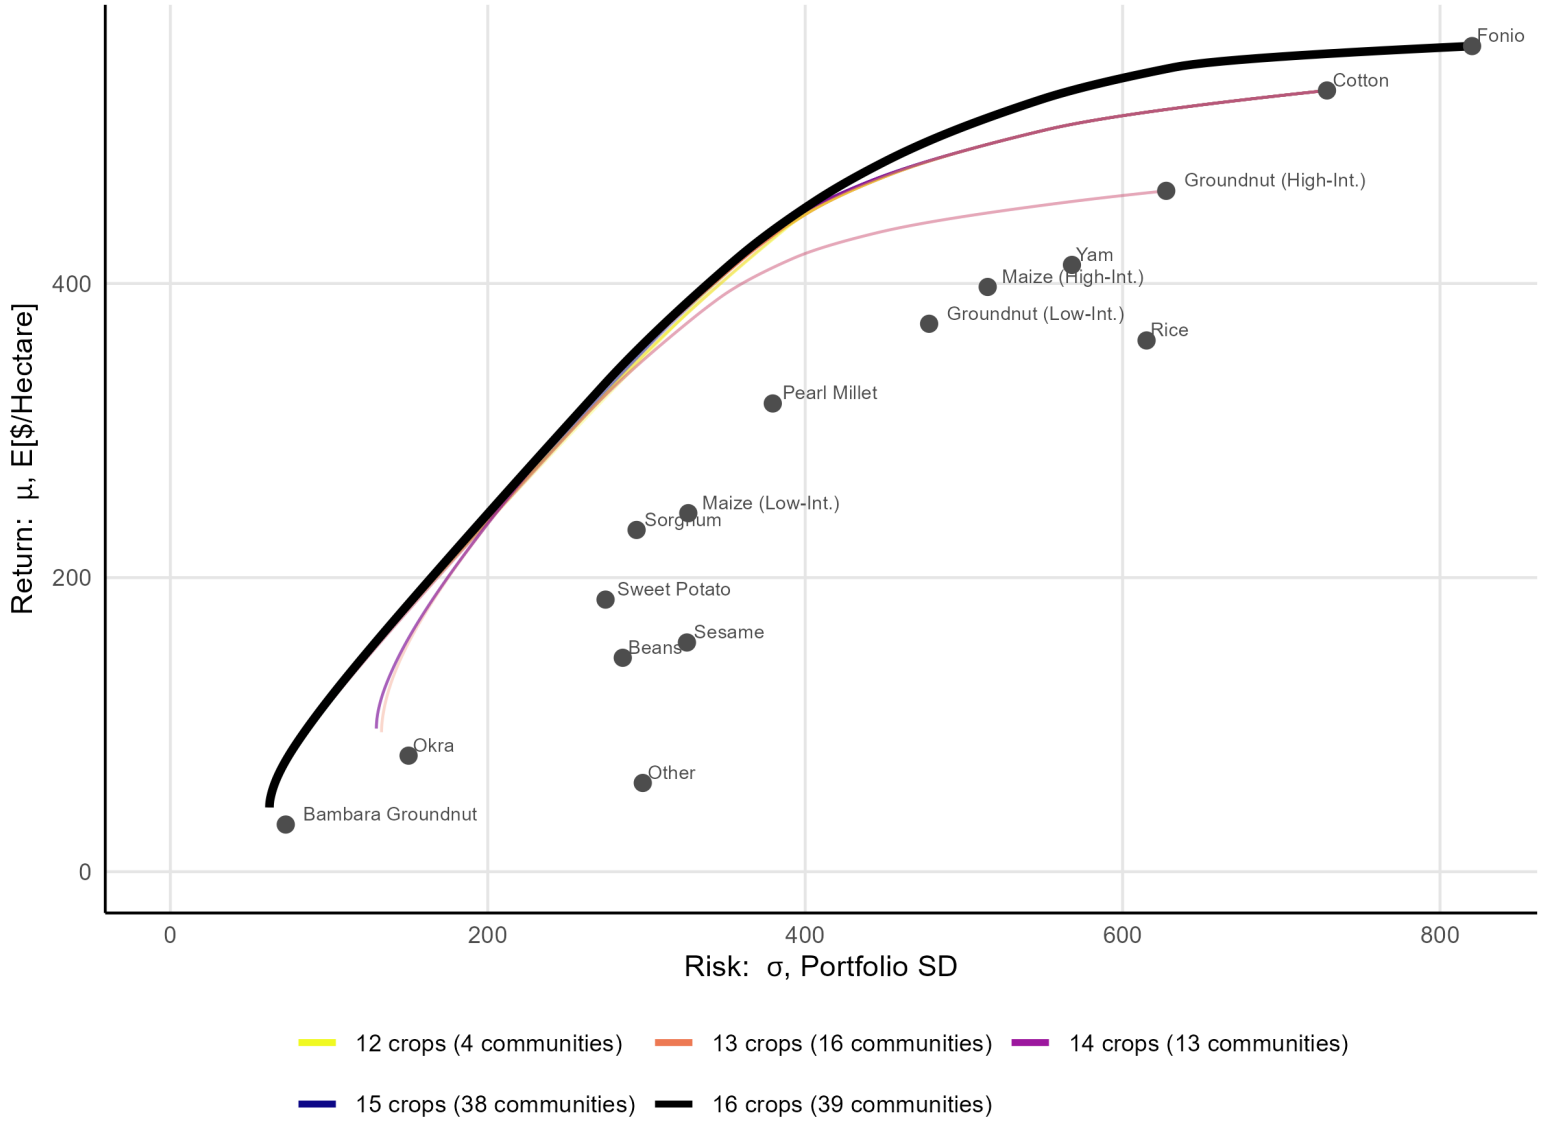
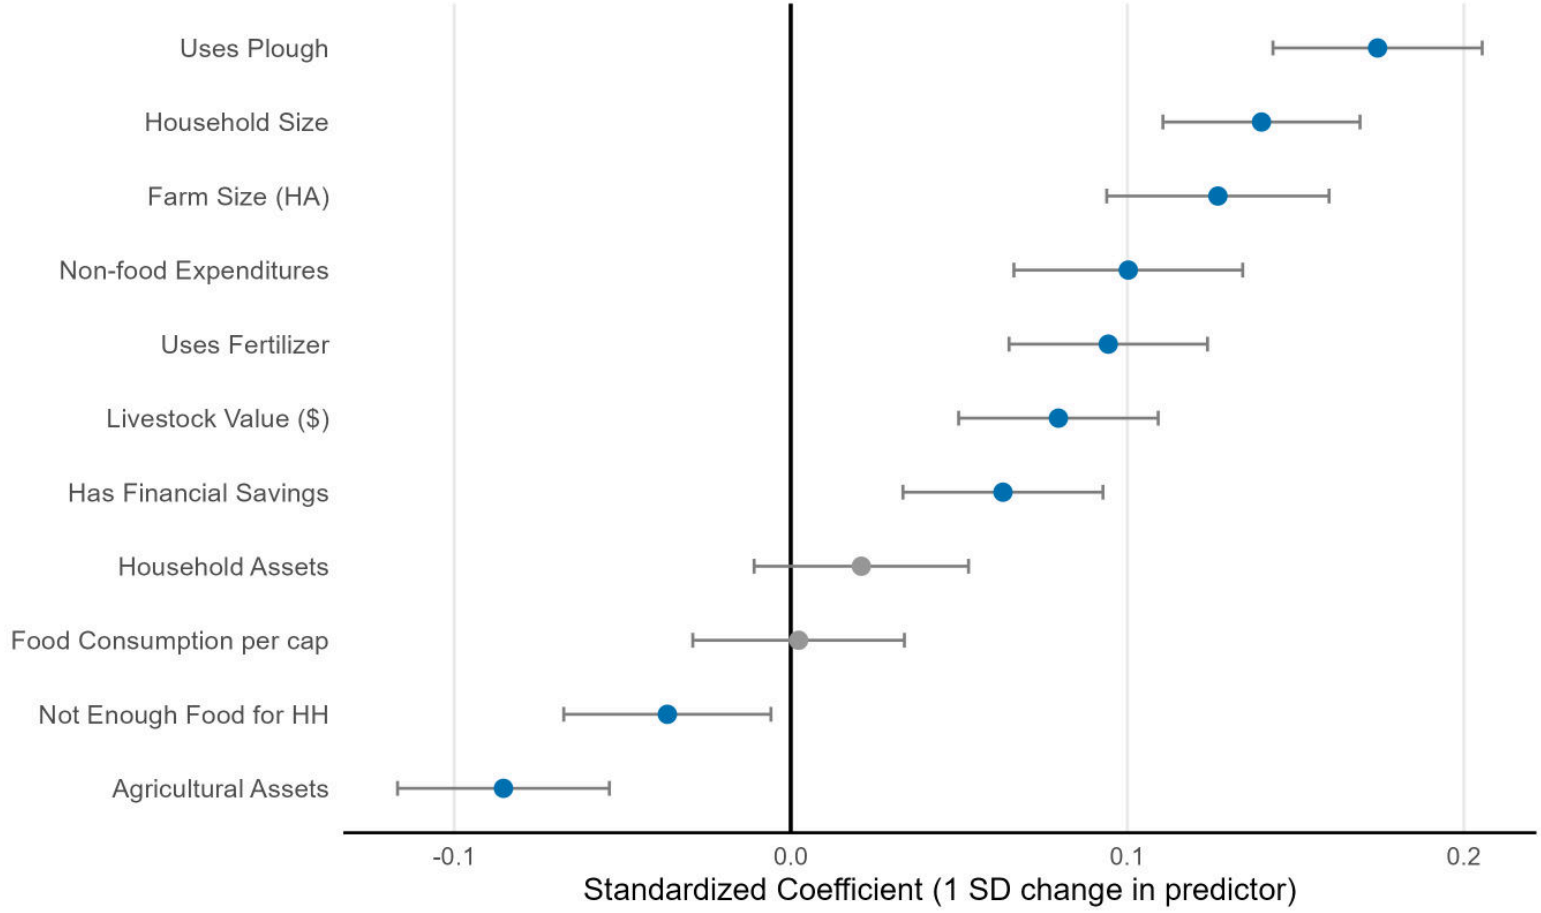
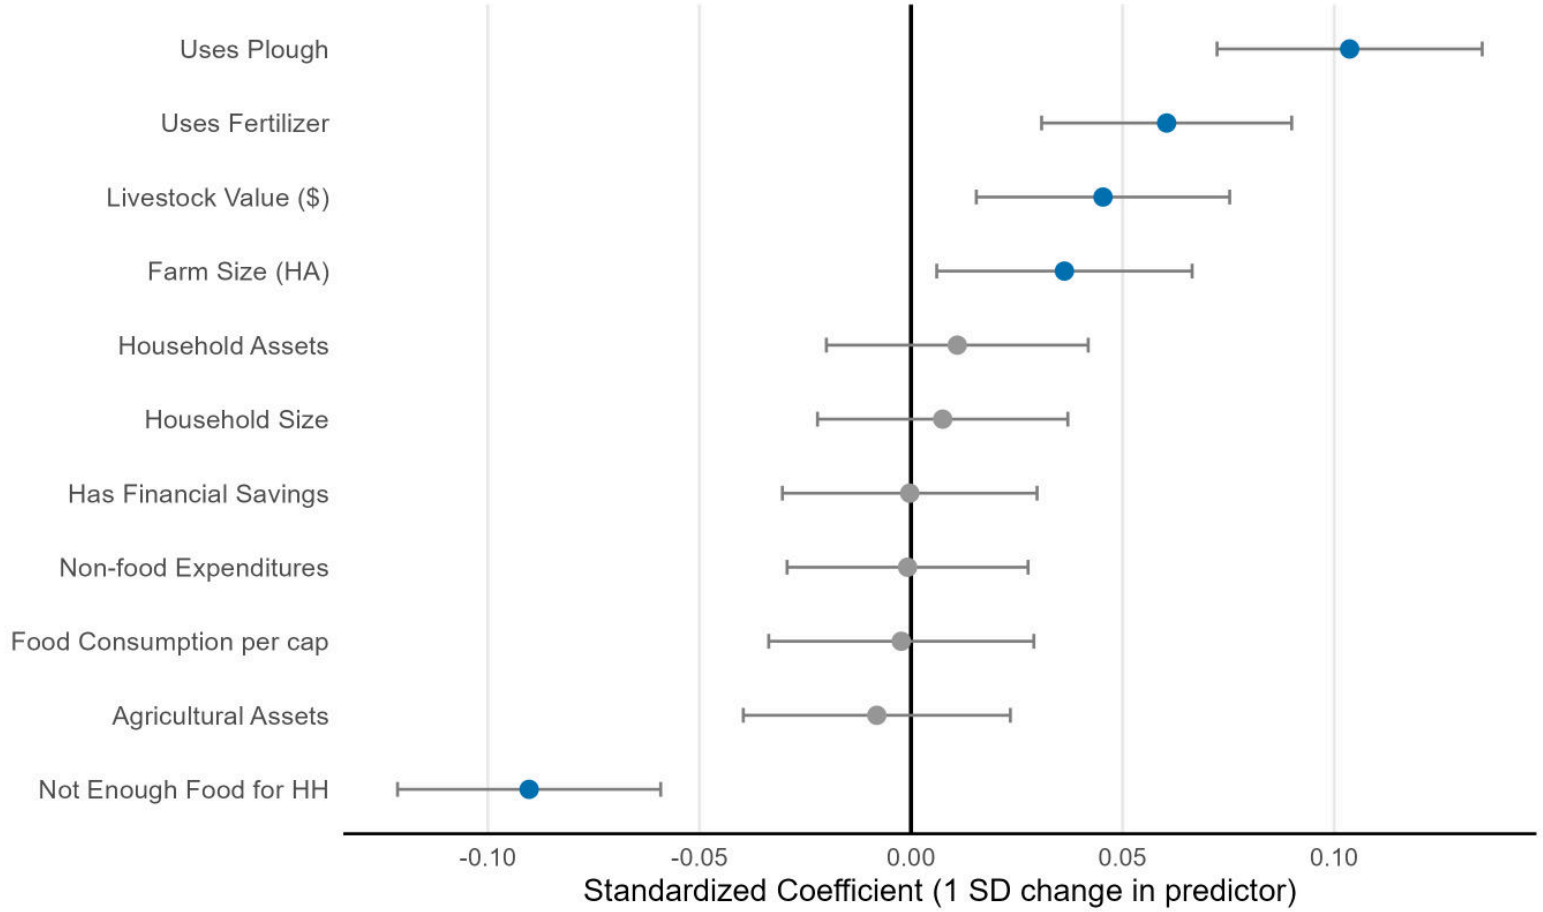


Figure A.16: What Predicts Portfolio Efficiency? Standardized Bivariate Associations



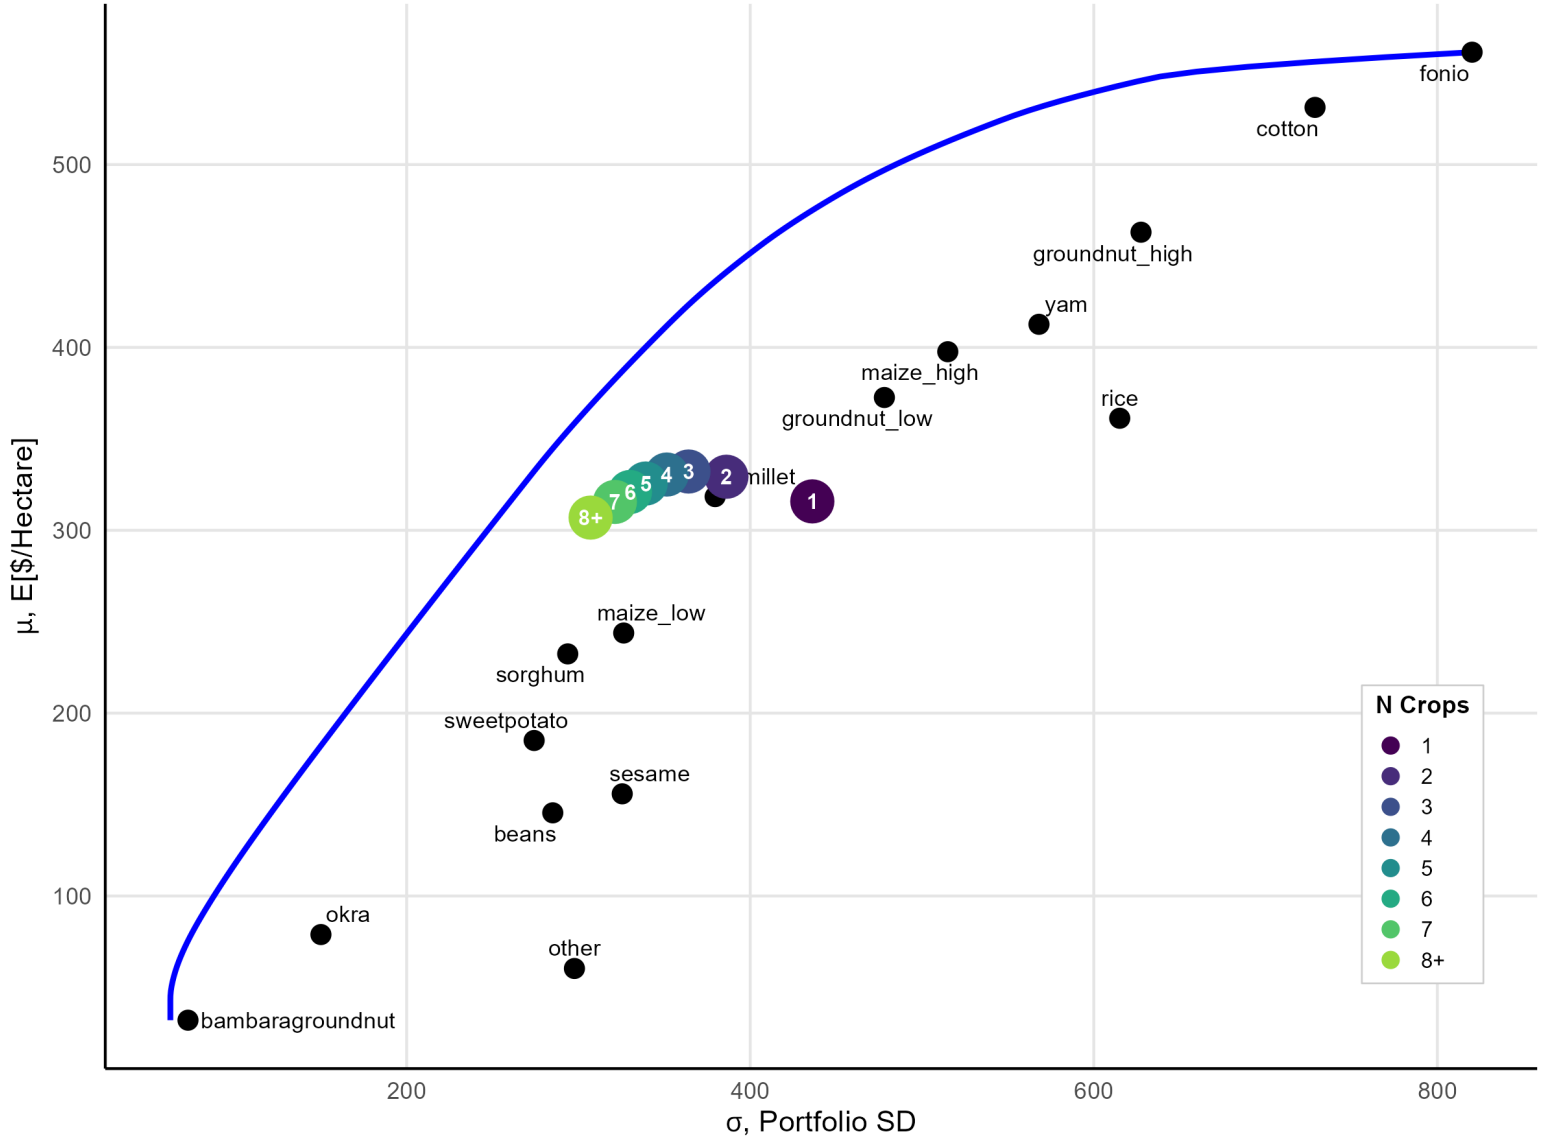
Notes: Each row is a separate bivariate regression of z(Efficiency) on z(predictor), where Efficiency is defined as the negative of the household's distance to its crop-set specific efficient frontier (so higher values = closer to frontier). Both LHS and RHS are standardized within the baseline (year 2010) sample. Crop-set fixed effects absorb the menu of crops available to the household, so the comparison is between households facing the same frontier. Standard errors are heteroskedasticity-robust. Baseline (2010) sample: 3,969 households across 18 crop sets and 110 communities. Bars show 95% confidence intervals and blue points indicate $p < 0.05$.

Figure A.17: What Predicts the Household's Risk Target? Standardized Bivariate Associations



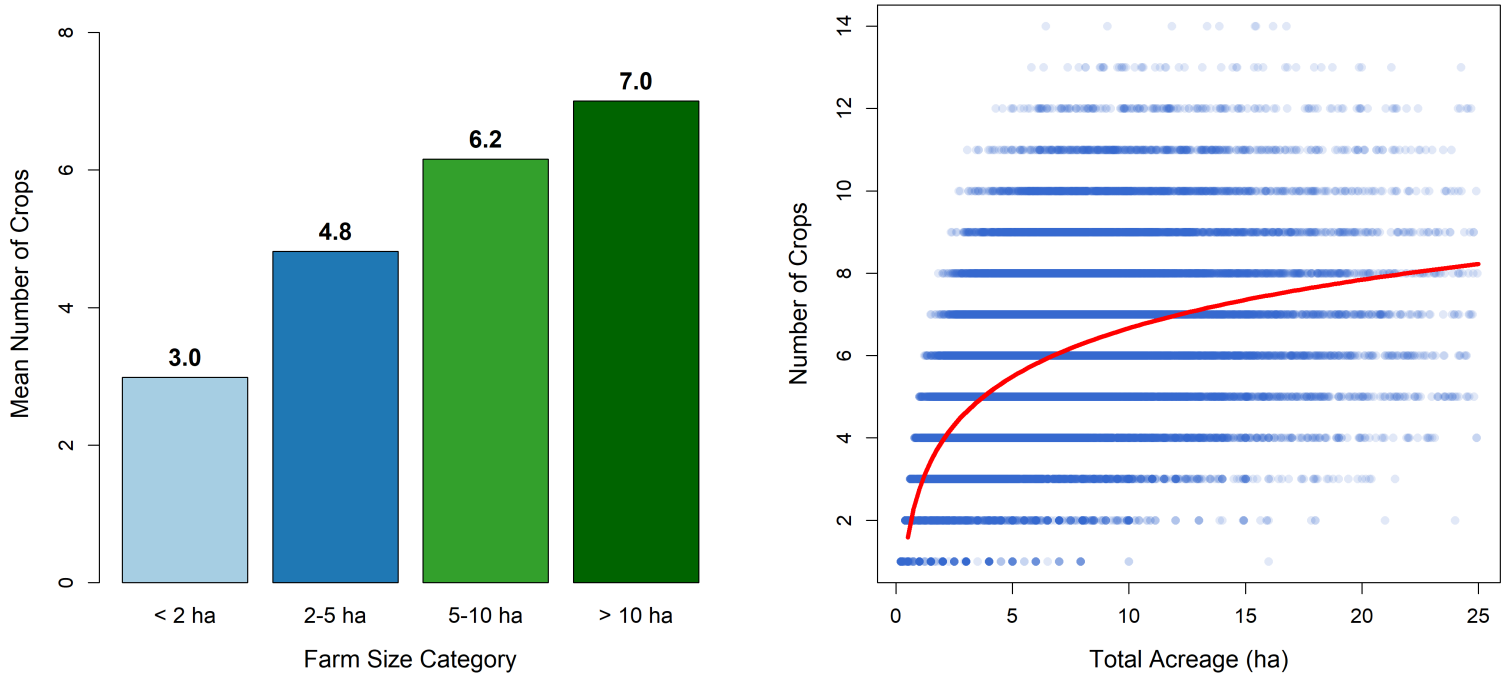
Notes: Each row is a separate bivariate regression of z(Frontier-Projected SD) on z(predictor). The frontier-projected standard deviation is σ at the point on the household's own crop-set specific efficient frontier closest to its observed portfolio: the level of risk the household appears to be targeting, as distinct from the level it realizes. Because every household sits inside its own frontier, the projected value is always below the realized one, and the difference is the displacement Figure A.16 explains. The projection is derived from the observed portfolio rather than measured independently, so these associations are descriptive rather than recovered preference parameters. Clustering instead by community leaves every verdict in this figure unchanged except farm size, which moves to $p = 0.076$; the reading in the text rests on the size of that coefficient relative to the constraint margin rather than on the threshold. Crop-set fixed effects absorb the menu of crops available to the household, so the comparison is between households facing the same frontier. Standard errors are heteroskedasticity-robust. Baseline (2010) sample: 3,969 households across 18 crop sets and 110 communities. Bars show 95% confidence intervals and blue points indicate $p < 0.05$.

Figure A.18: Mean Portfolio Position by Number of Crops Grown



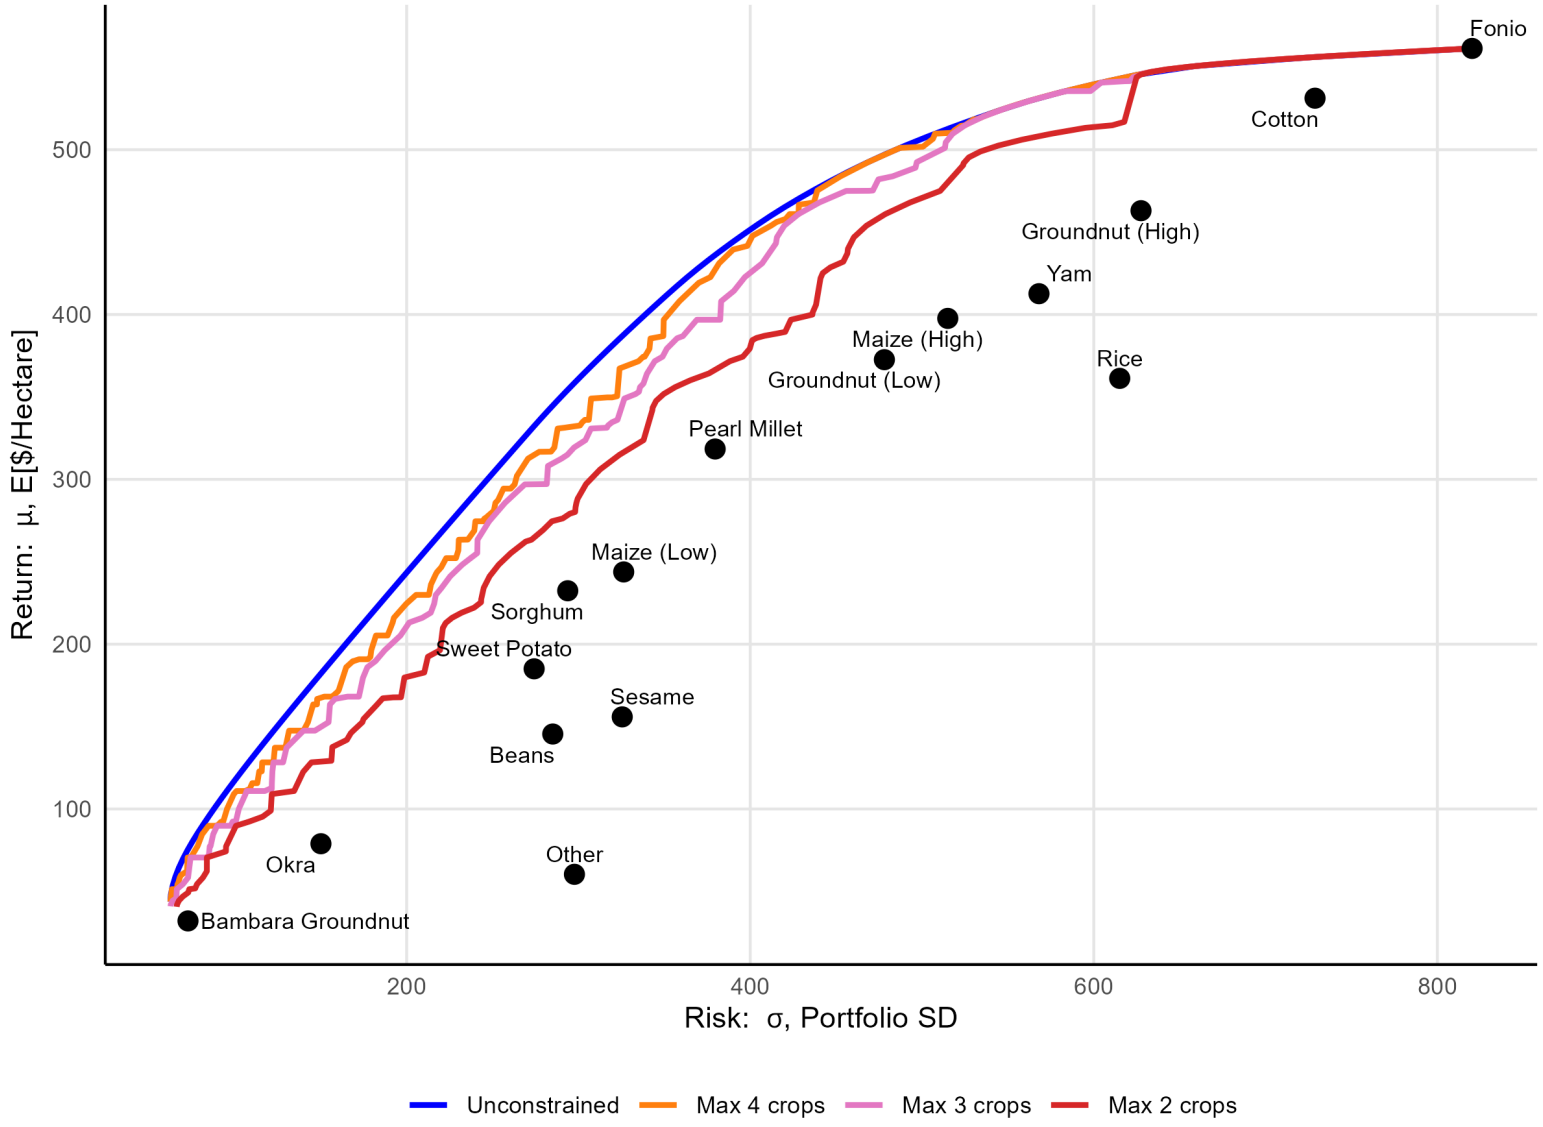
Notes: Each colored dot is the mean portfolio standard deviation and expected return of the household-years cultivating that number of crops (30,902 household-years, all arms and rounds; 439 at one crop, 5,568 at eight or more); black dots are the crop assets and the blue curve the full-menu efficient frontier of Figure A.14. Mean portfolio risk falls from $436 per hectare at one crop to $307 at eight or more, a 30% decline at nearly unchanged expected return.

Figure A.19: Size-diversification gradient



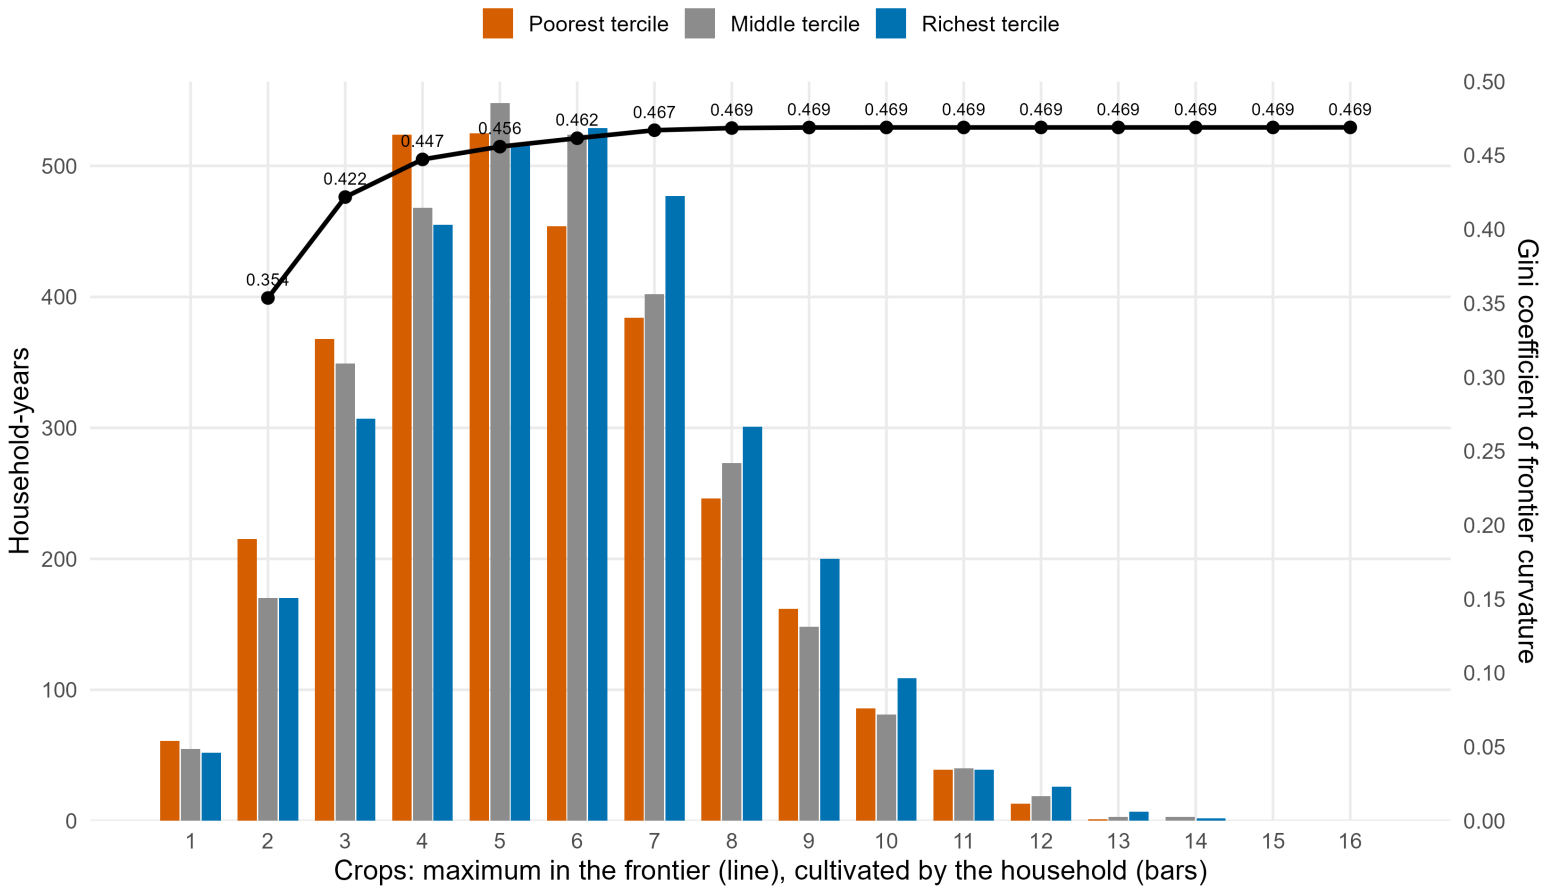
Notes: Number of distinct crop types cultivated, plotted against total cultivated area, at the household-year level. Doubling acreage is associated with $\hat{\beta}_{\log K} = 1.138$ additional crop types ($p < 0.001$, pre-floor cultivated area with community fixed effects), consistent with a per-crop fixed-cost mechanism that bounds cultivation cardinality from above by $\lfloor K/L_{c,\min} \rfloor$.

Figure A.20: Constrained efficient frontiers by maximum number of crops



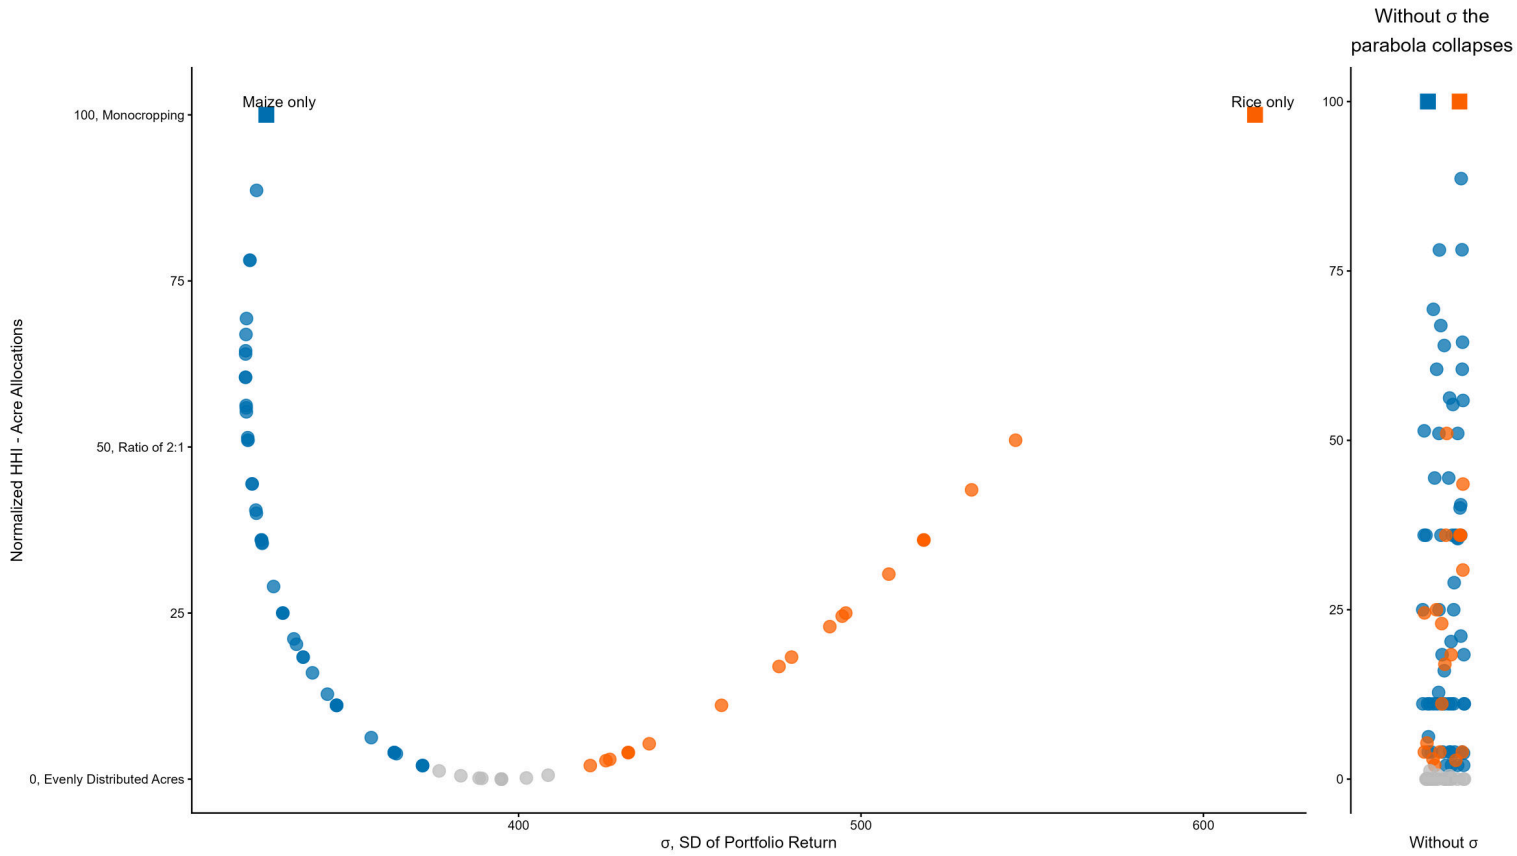
Notes: The light blue curve is the unconstrained efficient frontier (all 16 crops, no minimum-weight constraint), reproducing the frontier from Section 5. Each colored curve below it is a separate constrained efficient frontier with a fixed maximum cardinality (2, 3 or 4 crops). Each constrained frontier is itself the upper envelope of $\binom{16}{k}$ subset-specific frontiers, one per feasible k-crop combination drawn from the 16-crop menu.

Figure A.21: Diversification Room by Cardinality, against the Distribution of Crops Cultivated



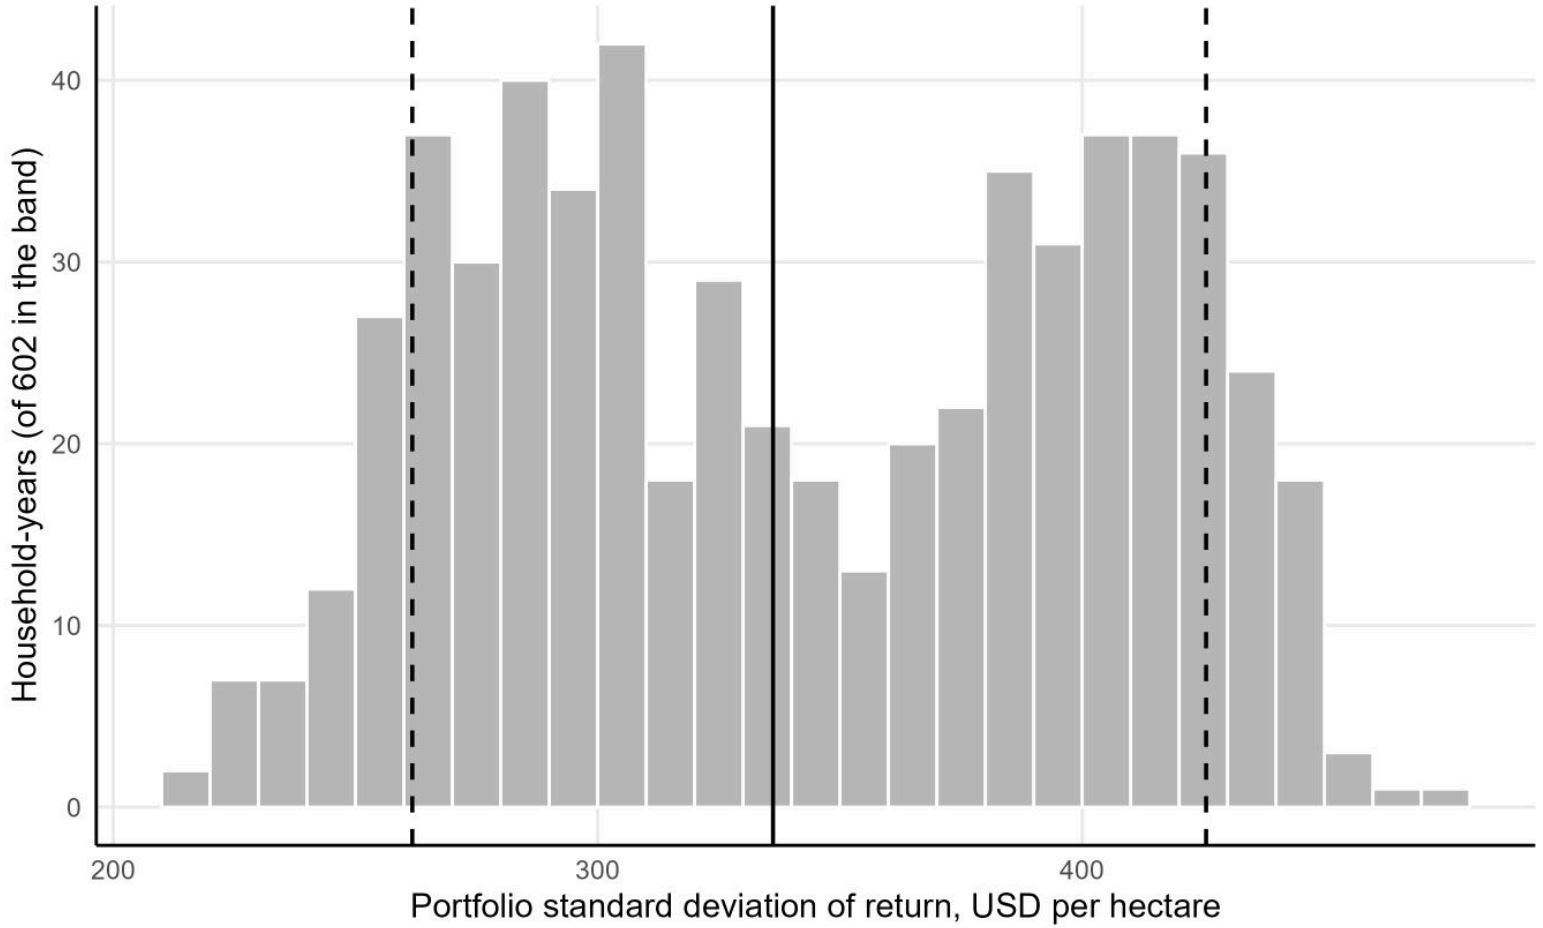
Notes: The line is the Gini coefficient of frontier curvature at each maximum cardinality K: the area between the K-crop constrained frontier and the straight chord joining the two corner crops (Bambara groundnut to fonio), divided by the area of the bounding triangle. It measures how far a K-crop menu bows the attainable set away from that chord, which is the diversification room the menu opens up. The K-crop frontier is the no-shorts frontier over every subset of at most K crops with no minimum-weight floor, built from 65,518 enumerated subsets; the elbow is the smallest K at which 90% of the rise from two crops to the full menu is in place. The bars are the distribution of unique crops cultivated per household-year among pure-control households across all four rounds (9,377 household-years, 2,394 households, the sample of Table B.4), split by baseline wealth tercile. Terciles are cut on the baseline wealth index exactly as in the heterogeneity analysis of Section 6.3, so a household in the richest tercile here is in the richest tercile there, and a household keeps its tercile in every round. The two are plotted against a common horizontal axis so that the cardinality at which a household sits can be read against the curvature its menu affords.

Figure A.22: Comparison of Different Measures used to Proxy for Diversification



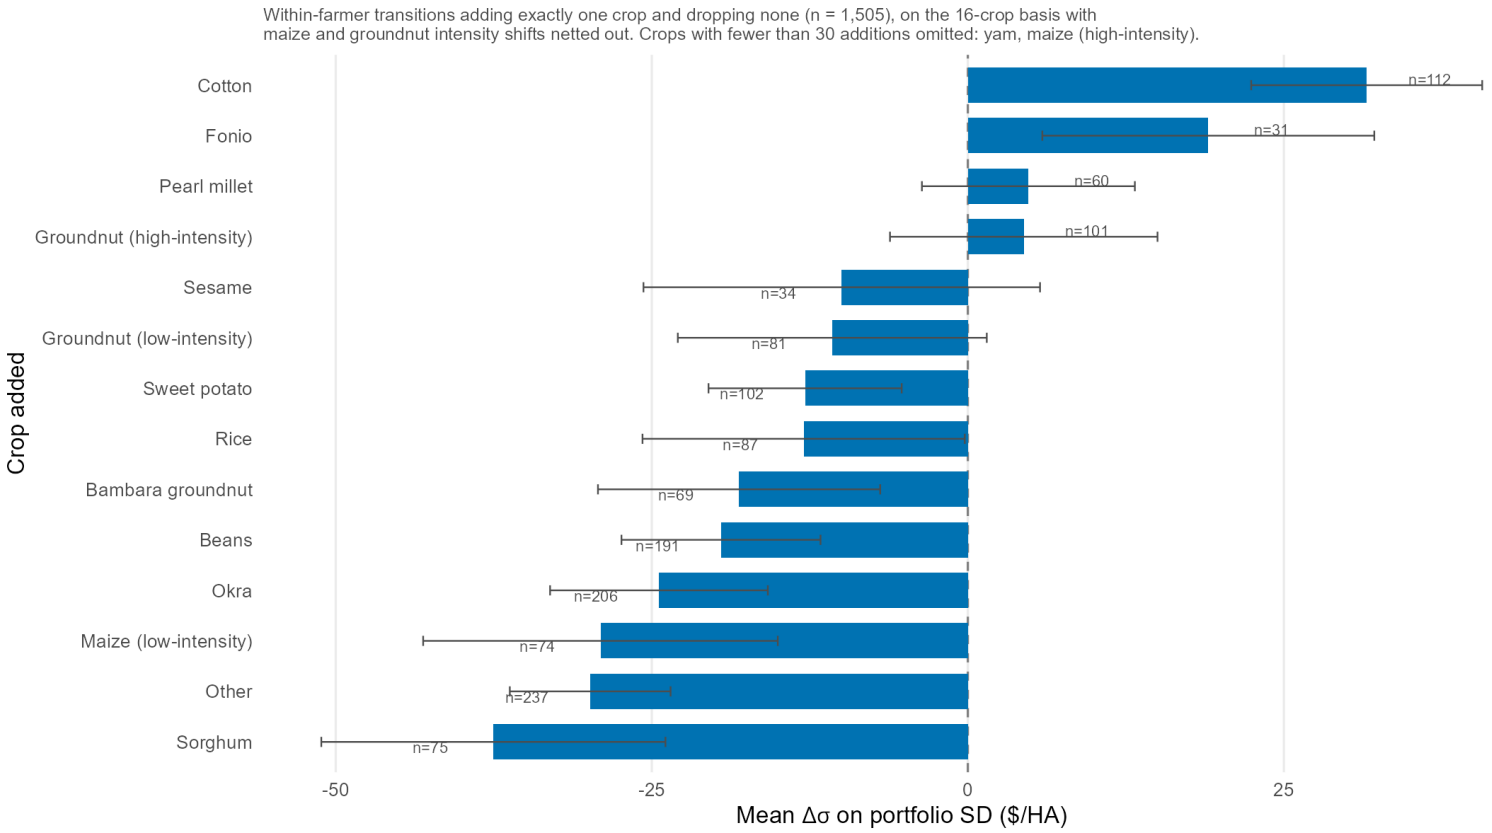
Notes: Household-years cultivating low intensity maize alone, rice alone, or maize and rice together, with no other crop. Maize here is the low-intensity variant of Section 2.2; household-years involving the high-intensity variant are excluded so that the two-crop portfolios trace a single parabola. The horizontal axis is the estimated portfolio standard deviation of return per hectare, the vertical axis the Normalized Herfindahl-Hirschman Index of land shares; the number line at right projects the same portfolios onto the NHHI alone.

Figure A.23: Portfolio Risk among Household-Years with 6 Crops and an NHHI between 15 and 20



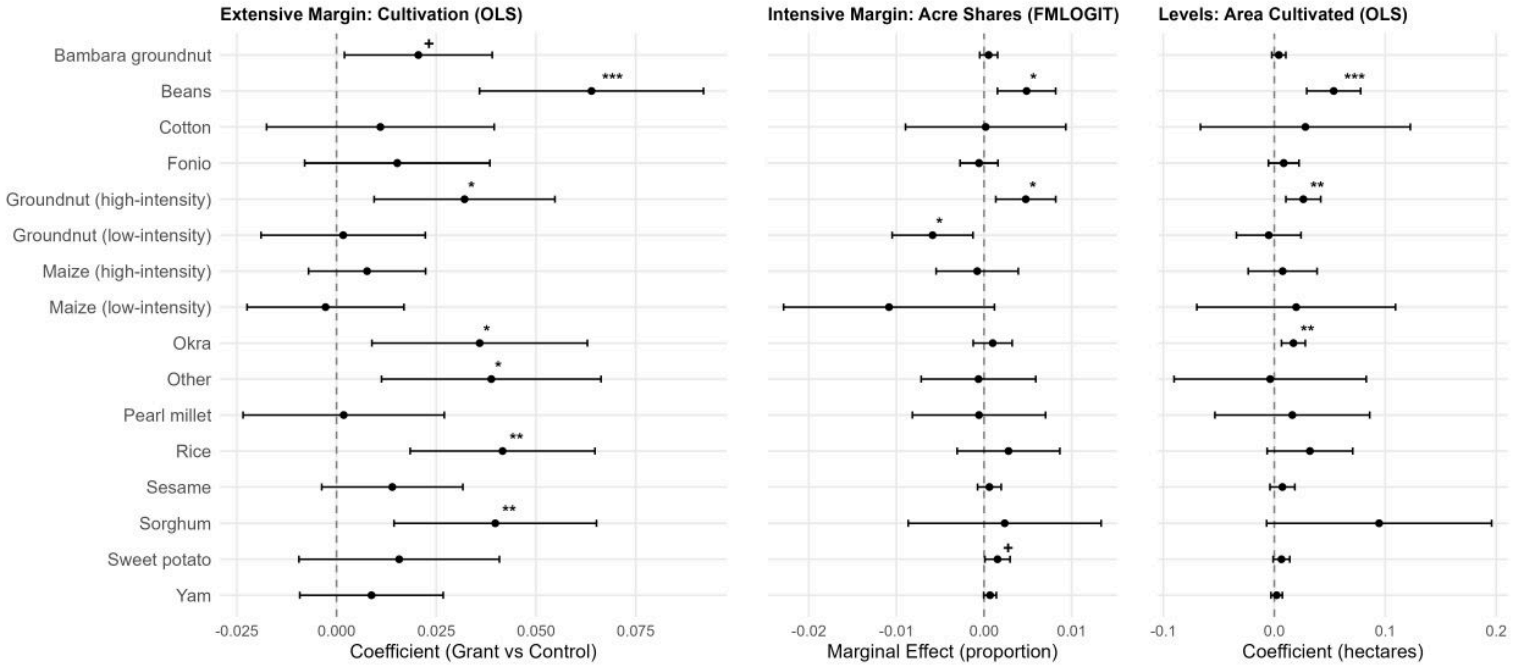
Notes: Distribution of portfolio standard deviation of return across the 602 household-years that cultivate exactly 6 crops with a Normalized Herfindahl-Hirschman Index between 15 and 20, so that every observation in the figure takes the same value of the crop count and, to within five points, the same value of the NHHI. Dashed lines mark the tenth and ninetieth percentiles ($262 and $426 per hectare), the solid line the median ($336); the full band runs from $211 to $475. The spread is not a thin tail: the distribution has two distinct modes, so households the count measures treat as identical divide into recognizably different risk positions. Sample is every household-year in the analysis panel with a positive portfolio standard deviation and both count measures defined, pooling all study arms and rounds.

Figure A.24: Within-farmer change in portfolio SD when one crop is added



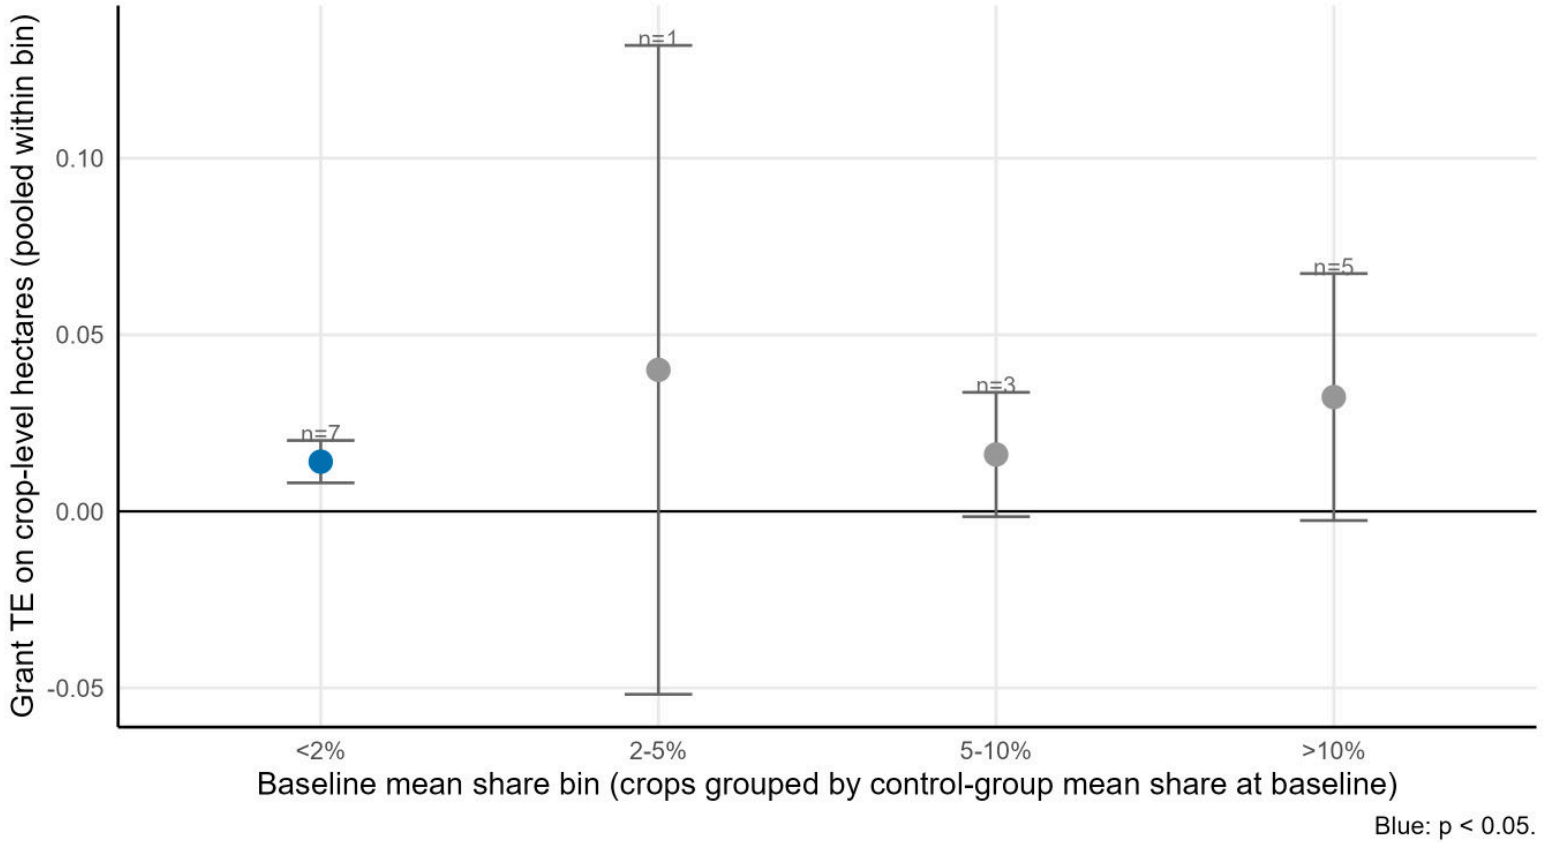
Notes: Sample restricted to within-farmer transitions between consecutive survey waves where the household added exactly one crop and dropped none ($n = 1{,}505$ transitions), on the 16-crop basis used throughout the paper. A move between the intensity variants of a single crop, low-intensity maize to high-intensity maize or the same for groundnut, is a substitution within that crop rather than an addition, so such a shift is netted out of both the added and the dropped set before the restriction is applied; a household that only intensifies is therefore excluded, and a household that intensifies while adding a crop is counted for the crop it added. For each transition, $\Delta\sigma$ is the change in portfolio standard deviation between the two waves. The bar for each crop reports the mean $\Delta\sigma$ across all transitions where that crop was the one added; error bars are 95% confidence intervals based on within-crop standard deviation. Crops with fewer than 30 such additions are omitted, 2 of them (yam, maize (high-intensity)). Cotton-adders move +\$31.6/HA on σ; okra-adders move −\$24.4/HA, a ∼\$56/HA wedge in σ-space at identical $\Delta N_{\text{crops}} = +1$. The widest gap on the figure is larger still: sorghum-adders move −\$37.5/HA.

Figure A.25: Treatment Effects by Crop: Extensive Margin (OLS), Intensive Margin (FMLOGIT), and Levels (OLS)



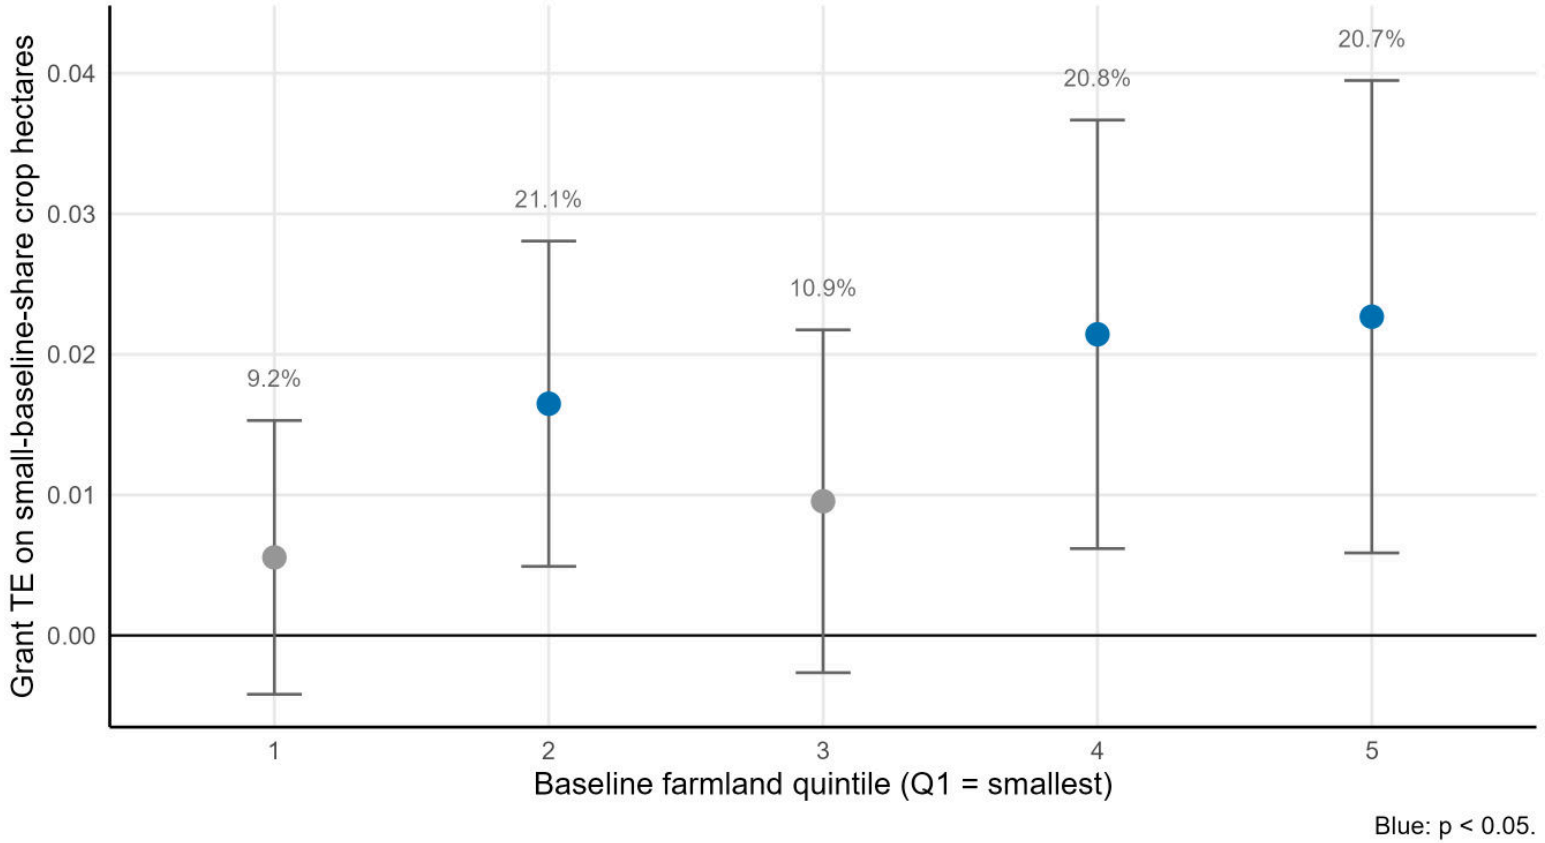
Notes: Left panel: OLS coefficients of grant receipt on the probability of cultivating each crop type (extensive margin). Middle panel: fractional multinomial logit (FMLOGIT) marginal effects of grant receipt on the share of cultivated land allocated to each crop (intensive margin). Right panel: OLS coefficients of grant receipt on hectares cultivated of each crop (levels). The two OLS panels follow the headline specification of Section 6: round indicators, the baseline value of the outcome, and the authors' preferred control suite. The FMLOGIT panel enters round indicators and a reduced control set (household structure and baseline farm size). Standard errors are clustered at the household level in all three panels. Sample restricted to grant-only villages (arm 4) versus pure control (arm 5), balanced panel. Bars show 90% confidence intervals. $+\ p < 0.1$, $*\ p < 0.05$, $**\ p < 0.01$, $***\ p < 0.001$.

Figure A.26: Hectare treatment effect by baseline-share bin



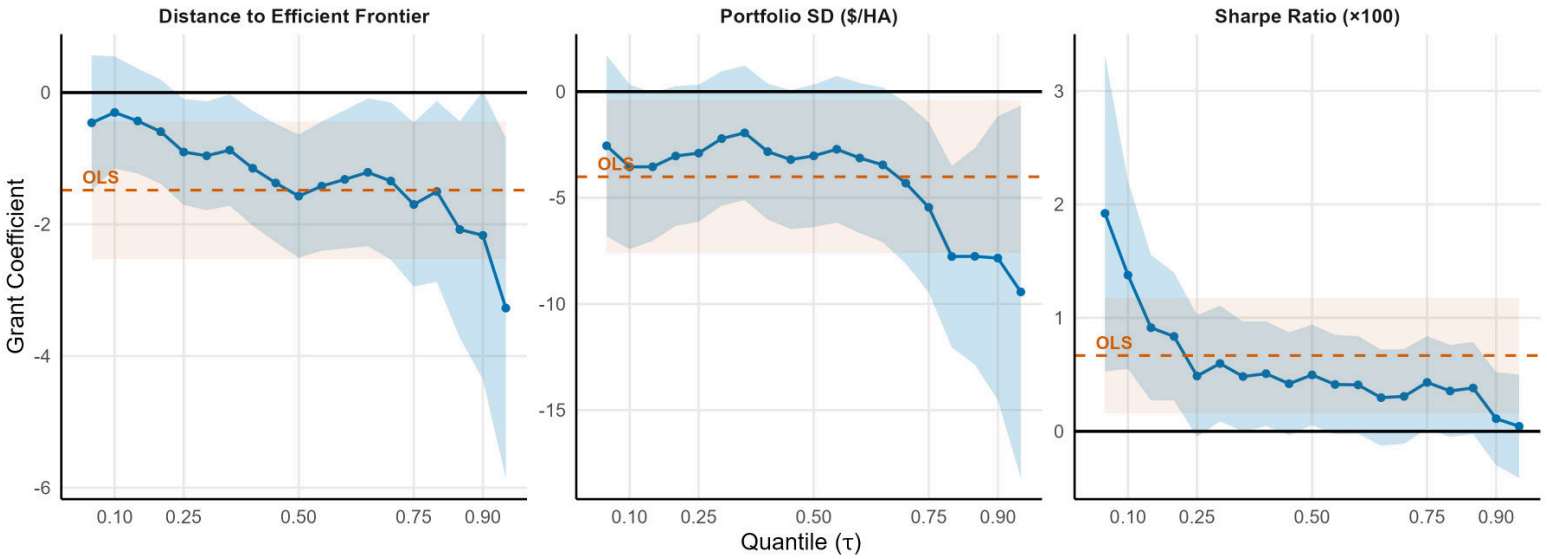
Notes: For each crop, baseline mean share is computed in the round-1 control group. Crops are then assigned to bins (0–2%, 2–5%, 5–10%, >10%), and the grant treatment effect on cultivated hectares is estimated pooled within each bin in a long farmer-crop-year panel with crop-set and crop fixed effects, controlling for baseline area and the authors' preferred control suite. Standard errors clustered at household. Labels report the hectare effect as a percentage of the bin's control-group mean.

Figure A.27: Extensive-margin treatment effect by baseline-farmland quintile



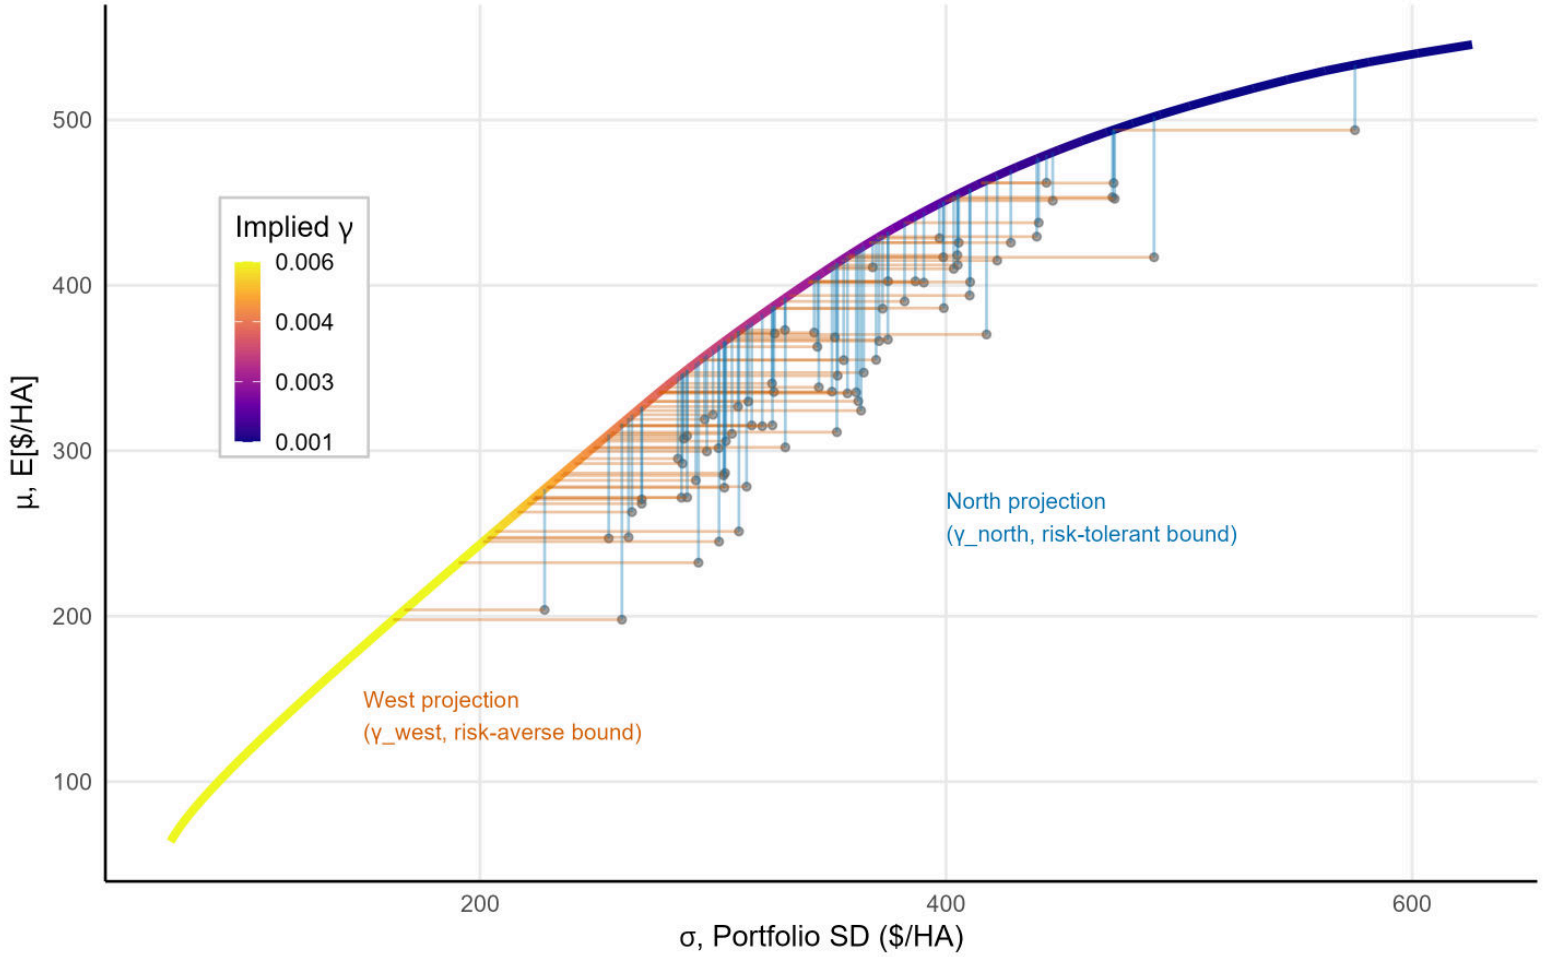
Notes: The grant treatment effect on cultivated hectares is estimated separately for each baseline-farmland quintile, pooled across the seven low-baseline-share crops (Bambara, beans, fonio, okra, sesame, sweet potato, yam). Specification matches Figure A.26: crop-set and crop fixed effects, baseline-area control, authors' preferred controls, SE clustered at household. Labels report the hectare effect as a percentage of the quintile's control-group mean.

Figure A.28: Quantile regression: grant effect across the outcome distribution



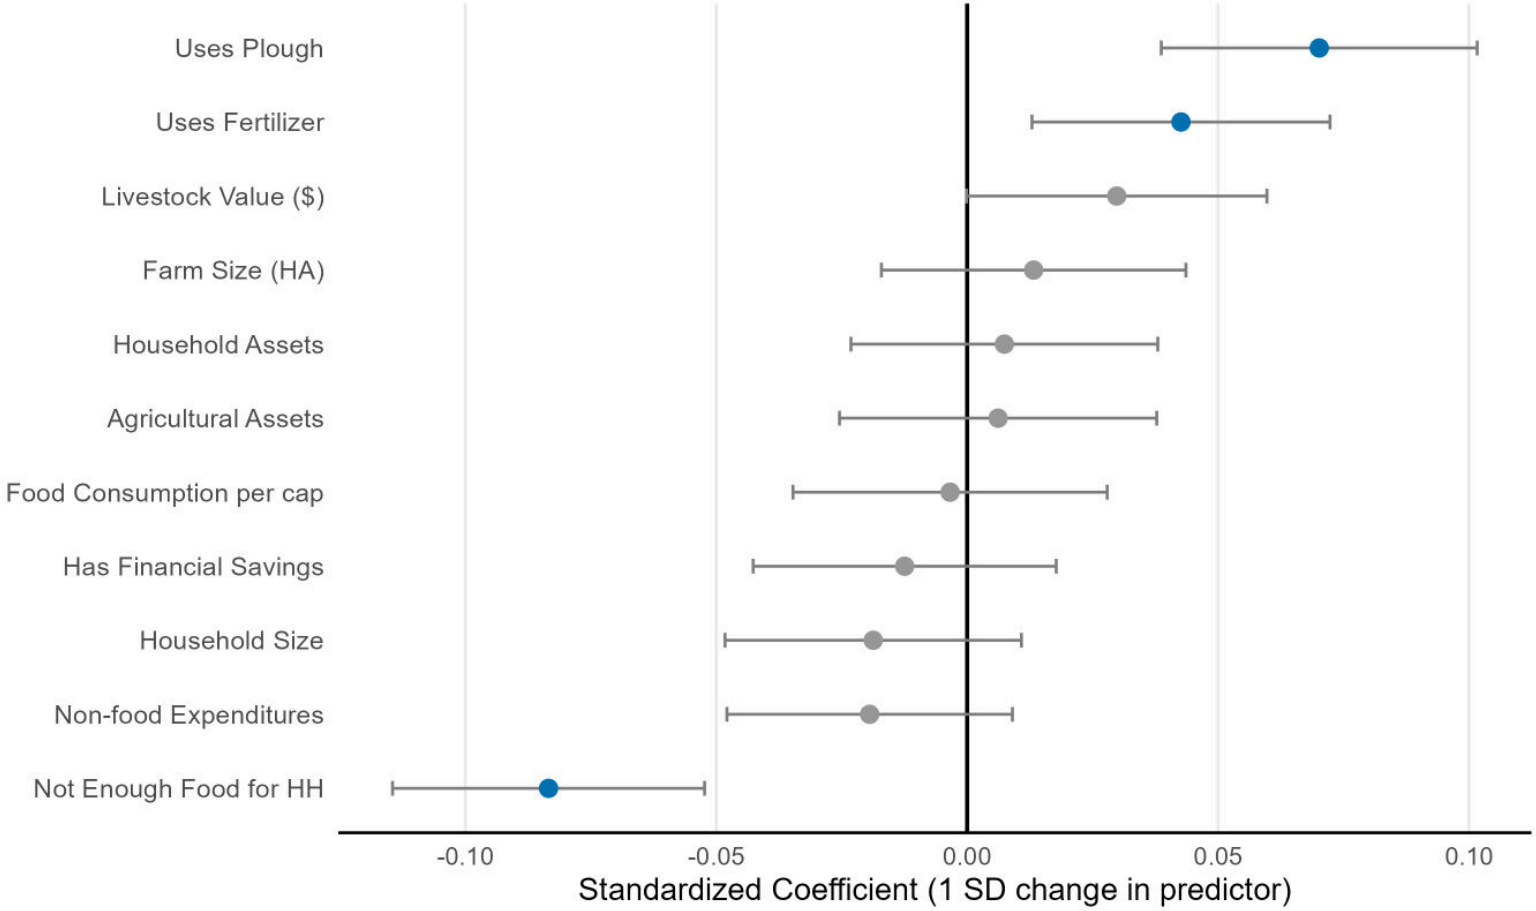
Notes: Quantile regression of three portfolio outcomes on grant assignment. Blue dots show QR coefficient estimates with 95% confidence intervals (kernel SE); the dashed orange line marks the OLS (mean) effect. For portfolio SD, the treatment effect grows from approximately −$3.5/HA at the 10th percentile to about −$7.8/HA at the 90th percentile, with the largest effects concentrated above the 70th percentile, a catch-up pattern in which high-σ households absorb the largest variance reductions. The same monotonic pattern appears for distance to the efficient frontier.

Figure A.29: Per-farmer frontier projections used to identify the γ bounds



Notes: Each gray dot is one baseline household's observed (σ_h, μ_h) position. Orange segments show the west projection (same μ, lower σ) used to identify γ_{west}; blue segments show the north projection (same σ, higher μ) used to identify γ_{north}. The frontier color encodes the implied γ at each frontier point: lower-σ portions of the frontier have steeper slope and therefore correspond to higher γ (more risk-averse).

Figure A.30: What Predicts Realized Portfolio Risk? Standardized Bivariate Associations



Notes: Each row is a separate bivariate regression of z(Portfolio SD) on z(predictor), using the household's realized portfolio standard deviation. This figure is the composite of the two margins the framework separates. Realized risk is the household's risk target plus the displacement its constraints impose, so a predictor that raises the target and simultaneously moves the household closer to the frontier has partly offsetting effects here. Figure A.17 reports the target margin and Figure A.16 the constraint margin; the coefficients below are attenuated toward zero relative to the target figure for that reason. Crop-set fixed effects absorb the menu of crops available to the household, so the comparison is between households facing the same frontier. Standard errors are heteroskedasticity-robust. Baseline (2010) sample: 3,969 households across 18 crop sets and 110 communities. Bars show 95% confidence intervals and blue points indicate $p < 0.05$. The lottery-based risk elicitations are excluded from the predictor set for the reliability reasons documented in Section 5.5.

Appendix B: Tables

Table B.1: Crop frequency across farming seasons, plot cultivation instances

Crop	2010	2011	2012	2017
Bambara groundnut	0	799	1991	360
Beans	1788	3811	3801	1460
Cotton	3399	3704	5383	4138
Fonio	841	1349	1431	430
Groundnut (high-intensity)	3074	2695	1670	3092
Groundnut (low-intensity)	5082	6501	7119	3614
Maize (high-intensity)	3561	1533	242	4179
Maize (low-intensity)	2363	5926	7652	1955
Okra	1075	2675	4035	1366
Other	2115	4195	4926	1873
Pearl millet	2685	2764	1977	1116
Rice	5173	6650	6731	3803
Sesame	285	593	987	440
Sorghum	5505	6402	5549	2182
Sweet potato	276	1678	2965	153
Yam	136	525	972	176

Notes: Number of plot-crop observations by crop type and survey round. Baseline (2010) precedes treatment; Follow-up 1 (2011) and Follow-up 2 (2012) are the post-treatment periods used in the headline specifications. The 2017 follow-up round contributes additional weather and price realizations to the four-wave covariance matrix estimation but is not used for treatment-effect estimation. Maize, sorghum, and millet are the most commonly cultivated crops across all periods.

Table B.2: Balance of Baseline Characteristics and Outcomes

Variable	Grant	Control	Difference	p	q_{BH}
Controls (8 variables; joint $F = 0.27$, $p = 0.97$)					
Extended household	0.699	0.711	-0.013	0.501	0.918
Polygamous	0.429	0.445	-0.016	0.417	0.894
Ag-assets (z)	0.631	0.613	0.018	0.726	0.960
HH-assets (z)	-0.043	-0.047	0.004	0.905	0.968
Food consumption per capita	1,827	1,821	6	0.887	0.968
Owns plough	0.596	0.579	0.017	0.390	0.894
Uses fertilizer	0.242	0.238	0.005	0.789	0.960
Farm size (HA)	0.503	0.495	0.008	0.619	0.947
Main outcomes (9 variables; joint $F = 2.00$, $p = 0.04$)					
Hectares cultivated	6.13	6.27	-0.14	0.403	0.894
Number of crop types	4.89	4.87	0.02	0.784	0.960
Normalized HHI (acres)	19.34	18.91	0.42	0.574	0.947
Total expenditure ($)	1,255	1,247	8	0.889	0.968
Total revenue ($)	4,747	4,845	-98	0.598	0.947
Total profit ($)	3,395	3,482	-86	0.495	0.918
Profit / HA ($)	547.1	549.2	-2.1	0.877	0.968
Expected profit / HA ($)	340.9	334.2	6.7***	0.009	0.129
Portfolio SD / HA ($)	356.3	347.2	9.2***	0.001	0.059
Efficiency outcomes (8 variables; joint $F = 2.34$, $p = 0.02$)					
Sharpe ratio	96.09	96.36	-0.27	0.416	0.894
Distance to crop-set frontier	47.80	46.89	0.92	0.261	0.825
Distance to crop-set k-frontier	33.81	32.71	1.10*	0.092	0.457
Sharpe at closest frontier point	117.1	117.7	-0.7***	0.003	0.079
Portfolio variance, diagonal	86,208	79,967	6,241***	0.005	0.083
Portfolio variance, off-diagonal	45,680	44,577	1,103*	0.100	0.457
Selection Ratio	169.1	172.5	-3.5**	0.016	0.172
Diversification Ratio	130.9	131.3	-0.3	0.476	0.918
Crop area shares (8 variables; joint $F = 2.43$, $p = 0.01$)					
Cotton	0.141	0.128	0.013*	0.077	0.425
Maize (low-intensity)	0.108	0.110	-0.003	0.744	0.960
Maize (high-intensity)	0.143	0.150	-0.007	0.380	0.894
Sorghum	0.215	0.224	-0.010	0.302	0.871
Pearl millet	0.104	0.110	-0.006	0.460	0.918
Rice	0.077	0.067	0.010**	0.037	0.287
Groundnut (low-intensity)	0.083	0.085	-0.002	0.658	0.960
Groundnut (high-intensity)	0.063	0.052	0.011**	0.030	0.275

Variable	Grant	Control	Difference	p	q_{BH}
Livestock counts (8 variables; joint $F = 0.14$, $p = 1.00$)					
Poultry	5.56	5.69	-0.13	0.764	0.960
Goats	1.45	1.45	0.00	0.994	0.994
Sheep	1.19	1.20	-0.01	0.915	0.968
Calves	0.490	0.516	-0.026	0.723	0.960
Cows	1.47	1.49	-0.02	0.913	0.968
Donkeys	0.250	0.264	-0.014	0.620	0.947
Oxen	0.346	0.327	0.019	0.796	0.960
Traction oxen	0.724	0.710	0.014	0.803	0.960
Non-crop assets and enterprise (3 variables; joint $F = 0.15$, $p = 0.93$)					
Livestock value (FCFA)	460,553	459,862	691	0.987	0.994
Has a business	0.549	0.536	0.013	0.528	0.937
Enterprise profit, 12 months (FCFA)	69,294	69,205	90	0.992	0.994
Input expenditure (3 variables; joint $F = 2.01$, $p = 0.11$)					
Fertilizer expenditure (FCFA)	24,040	19,182	4,858*	0.065	0.425
Manure expenditure (FCFA)	113.1	176.6	-63.4	0.317	0.871
Other input expenditure (FCFA)	10,913	9,473	1,440	0.126	0.534
Demographics (8 variables; joint $F = 1.07$, $p = 0.38$)					
Household size	7.40	7.48	-0.08	0.578	0.947
Number of adults	3.30	3.25	0.05	0.423	0.894
Number of children	4.08	4.20	-0.13	0.228	0.825
Number of working age	3.60	3.57	0.03	0.724	0.960
Average age	21.75	21.17	0.58	0.145	0.570
Average adult age	36.38	35.96	0.41	0.270	0.825
Share of adults male	0.468	0.456	0.011*	0.075	0.425
Leadership position in household	0.282	0.262	0.020	0.270	0.825

Notes: Grant (arm 4, $N = 800$) and pure-control (arm 5, $N = 2,381$) households at baseline (2010), authors' preferred sample. Difference is the coefficient on the grant indicator from a regression of each variable on it, with heteroskedasticity-robust standard errors; p is its p-value and q_{BH} the Benjamini-Hochberg adjusted value across all 55 tests. Group headers report the classical joint F from a regression of the grant indicator on the group's variables; across all 55 variables together, $F(55, 3119) = 0.93$, $p = 0.62$. Outcome variables are defined in Sections 2.6 and 6; livestock, enterprise, input and demographic variables are from the household roster, in FCFA where monetary. Stars refer to the unadjusted p: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.3: Summary Statistics: Survey-Measured Outcomes

	N	Mean	SD	Min	p25	Median	p75	Max
Hectares Cultivated	9,377	7.02	4.57	0.20	3.94	6.02	9.01	55.42
N Crop Types	9,377	5.56	2.25	1.00	4.00	5.00	7.00	14.00
Normalized HHI	9,377	18.40	17.32	0.00	8.21	13.78	22.25	100.00
Total Expenditure (USD)	9,377	1,079.71	1,426.62	0.00	254.99	643.77	1,412.15	36,651.68
Total Revenue (USD)	9,377	3,464.62	6,180.71	0.00	894.44	1,974.58	4,283.37	283,153.52
Total Profit (USD)	9,377	2,282.71	2,800.79	-1,367.64	509.67	1,292.54	2,995.64	27,647.63
Profit/HA (USD)	9,377	315.38	306.88	-578.83	102.84	217.92	465.03	4,244.00

Notes: Pooled summary statistics across all rounds (baseline through 2017) for households in pure control villages (arm 5), balanced panel. Treated households are excluded so that the table describes the setting rather than the post-treatment distribution. All monetary values in 2017 USD.

Table B.4: Summary Statistics: MPT-Constructed Portfolio Outcomes

	N	Mean	SD	Min	p25	Median	p75	Max
Expected Profit/HA (USD)	9,377	325.46	68.81	60.36	275.57	321.84	375.09	561.45
Portfolio SD (USD)	9,377	342.06	70.29	174.25	290.14	327.34	385.85	820.22
Sharpe Ratio (×100)	9,377	95.25	8.91	20.28	91.10	96.75	101.24	114.07
Euclidean Dist to Crop-Set EF	9,377	48.50	19.64	0.00	35.64	45.35	57.47	191.35
Sharpe at Ideal Point	9,377	117.79	5.33	68.45	116.71	119.92	121.24	121.71
Selection Ratio	9,377	171.20	34.43	100.00	144.17	164.47	195.24	287.17
Diversification Ratio	9,377	133.57	11.94	100.00	126.47	134.30	141.43	170.95

Notes: Pooled summary statistics across all rounds (baseline through 2017) for households in pure control villages (arm 5), balanced panel. Treated households are excluded so that the table describes the setting rather than the post-treatment distribution. Portfolio outcomes constructed using crop return moments estimated from pure control households. All monetary values in 2017 USD.

Table B.5: Crop Reference

Crop	Survey term	Type	Description
Rice	Riz	Cereal	Cereal grown in lowland and seasonally flooded fields.
Pearl millet	Petit mil	Cereal	Cereal grown on upland fields; a staple grain of the region.
Sorghum	Sorgho	Cereal	Cereal grown on upland fields.
Maize	Maïs	Cereal	Cereal grown for grain; split into high- and low-intensity variants in this paper (Section 2.2).
Fonio	Fonio	Cereal	Small-grained cereal with a short growing season.
Sesame	Sésame	Oilseed	Oilseed grown mainly for sale.
Cotton	Coton	Fibre	Fiber crop grown for sale.
Groundnut	Arachide	Legume	Legume grown for oil, food and sale; split into high- and low-intensity variants (Section 2.2).
Beans	Haricot	Legume	Grain legume; the roster's single bean category, covering cowpea and other beans.
Okra	Gombo	Vegetable	Vegetable grown for its pods.
Yam	Igname	Tuber	Starchy tuber.
Sweet potato	Patate douce	Tuber	Starchy tuber with a short cycle.
Bambara groundnut	Pois de terre	Legume	Legume with underground pods; unfortunately absent from the baseline roster (Appendix C.6).
Other	—	Residual	Residual category pooling crops too infrequent to model separately.

Notes: The survey term is the label used on the enumerator's crop roster. Groundnut and maize each enter the analysis as two assets, a high-intensity and a low-intensity variant, following the cost-per-hectare split described in Section 2.2; the split distinguishes input intensity within a species rather than different plants, so each appears once here. "Other" is an accounting residual rather than a crop: it pools species cultivated too rarely to support crop-level moment estimation. Descriptions are intended only to orient the reader to what kind of plant each asset is.

Table B.6: Estimated Moments of Crop Returns (mean, standard deviation, skewness, excess kurtosis)

Crop	E[$/HA]	SD	Skew	Ex.Kurt
Fonio	561	820	2.22	12.34
Cotton	531	729	3.70	71.88
Groundnut (high-intensity)	463	627	1.18	5.12
Yam	413	568	0.38	7.36
Maize (high-intensity)	398	515	0.42	5.96
Groundnut (low-intensity)	373	478	2.27	13.12
Rice	361	615	2.27	11.58
Pearl millet	318	380	1.54	10.98
Maize (low-intensity)	244	326	1.88	17.73
Sorghum	232	294	1.73	9.17
Sweet potato	185	274	1.03	7.67
Sesame	156	325	2.08	16.69
Beans	145	285	2.26	20.47
Okra	79	150	1.78	9.48
Other	60	298	4.61	48.78
Bambara groundnut	32	73	0.66	7.93

Notes: Expected return (μ) is the average profit per hectare in 2017 USD, estimated from farmer-level averages across seasons. Standard deviation (σ) is estimated from squared deviations of farmer-crop-year observations from farmer-crop means. Skewness and excess kurtosis are computed from the same farmer-demeaned residuals: skewness as $E[(R-\mu)^3]/\sigma^3$ and excess kurtosis as $E[(R-\mu)^4]/\sigma^4 - 3$, where excess kurtosis equals zero under a Gaussian return distribution. Sample restricted to farmers cultivating each crop in at least two seasons to ensure time variation in the estimates.

Table B.7: Estimated Correlation Matrix of Crop Profits per Hectare

	Bam	Bea	Cot	Fon	Gro_H	Gro_L	Mai_H	Mai_L	Okr	Oth	Pea	Ric	Ses	Sor	Swe	Yam
Bambara groundnut	1.00															
Beans	0.40	1.00														
Cotton	0.08	0.27	1.00													
Fonio	0.21	0.41	0.32	1.00												
Groundnut (High)	0.23	0.36	0.22	0.38	1.00											
Groundnut (Low)	0.27	0.42	0.29	0.51	0.34	1.00										
Maize (High)	-0.02	0.30	0.19	0.38	0.23	0.31	1.00									
Maize (Low)	0.15	0.25	0.30	0.32	0.33	0.38	0.40	1.00								
Okra	0.03	0.18	0.26	0.29	0.33	0.42	0.09	0.23	1.00							
Other	-0.18	0.13	0.00	0.29	0.09	0.24	0.14	0.07	0.12	1.00						
Pearl millet	0.23	0.37	0.39	0.47	0.36	0.46	0.28	0.41	0.23	0.16	1.00					
Rice	0.19	0.37	0.28	0.51	0.38	0.47	0.26	0.35	0.33	0.15	0.44	1.00				
Sesame	0.17	0.22	0.17	0.29	0.33	0.28	0.20	0.19	0.07	0.19	0.31	0.31	1.00			
Sorghum	0.26	0.46	0.37	0.50	0.42	0.50	0.39	0.46	0.27	0.22	0.56	0.43	0.33	1.00		
Sweet potato	0.28	0.38	0.23	0.35	0.33	0.35	0.13	0.11	0.45	0.27	0.21	0.26	0.44	0.27	1.00	
Yam	0.44	0.49	0.29	0.49	0.31	0.32	0.41	0.13	0.37	0.11	0.28	0.31	0.34	0.37	0.41	1.00

Notes: Lower triangular correlation matrix of crop profits per hectare. Correlation coefficients estimated from farmer-level profit deviations across crop pairs. Sample restricted to farmers cultivating both crops in at least two seasons. Cell shading indicates correlation sign and magnitude: red cells denote negative correlations (hedging opportunities), blue cells denote positive correlations. Only 2 of the 120 crop pairs correlate negatively, Bambara groundnut with the residual "other" category (−0.18) and Bambara groundnut with high-intensity maize (−0.02), so the diversification available in this setting does not come from true offsets. It comes from the spread in positive comovement: pairwise correlations run from essentially zero (cotton and "other", +0.001) to +0.56 (pearl millet and sorghum), around a mean of +0.29. A household combining crops at the low end of that range holds substantially less risk than one combining crops at the high end, even though every pair moves together on average. Table B.8 reports how many farmer-year observations stand behind each entry, so the thinner pairs can be weighed accordingly. Column headers show three-letter crop abbreviations.

Table B.8: Sample Sizes Behind Each Entry of the Covariance Matrix

	Bam	Bea	Cot	Fon	Gro_H	Gro_L	Mai_H	Mai_L	Okr	Oth	Pea	Ric	Ses	Sor	Swe	Yam
Bambara groundnut	**416**															
Beans	135	**3,953**														
Cotton	157	1,360	**5,017**													
Fonio	60	353	617	**1,427**												
Groundnut (High)	72	552	1,157	285	**3,256**											
Groundnut (Low)	232	1,975	3,710	832	1,101	**10,831**										
Maize (High)	39	442	1,289	235	744	1,363	**2,136**									
Maize (Low)	182	1,323	2,650	589	1,134	3,637	121	**5,637**								
Okra	88	474	945	191	687	1,473	402	1,237	**2,880**							
Other	143	1,230	1,800	660	816	2,448	640	1,725	698	**5,849**						
Pearl millet	81	890	1,174	307	329	1,714	388	1,105	213	790	**2,283**					
Rice	218	2,004	3,621	851	1,586	4,808	1,341	3,397	1,136	2,501	1,765	**10,567**				
Sesame	20	159	234	67	125	306	80	235	83	191	125	330	**505**			
Sorghum	189	1,920	2,785	652	1,036	3,992	918	2,879	949	1,948	1,595	4,087	250	**6,572**		
Sweet potato	60	472	512	136	160	714	120	520	178	474	325	742	92	652	**899**	
Yam	28	196	281	96	79	321	60	243	82	230	130	343	32	286	158	**366**

Notes: Sample sizes behind each entry of Tables B.7 and B.9. Bold diagonal entries count the plot-crop-year deviation observations behind each crop's variance, running from 366 to 10,831, of which 214 to 8,782 survive as degrees of freedom once the $N - G$ correction absorbs the farmer-crop means; a household contributing several plot-crop instances of the same crop in one season contributes each of them here. Off-diagonal entries count the farmer-year observations in which a household cultivated both crops, the overlap of two cultivator sets rather than the sample size, which is the unit the correlations are estimated on. These run from 20 (Bambara groundnut with sesame) to 4,808 with a median of 474, and 17 of the 120 pairs rest on fewer than one hundred farmer-years. Cells are shaded relative to the off-diagonal median, red for thinner and blue for thicker. Column headers show three-letter crop abbreviations.

Table B.9: Estimated Variance-Covariance Matrix of Crop Profits per Hectare (squared USD/HA)

	Bam	Bea	Cot	Fon	Gro_H	Gro_L	Mai_H	Mai_L	Okr	Oth	Pea	Ric	Ses	Sor	Swe	Yam
Bambara groundnut	**5,287.8**															
Beans	8,247.9	**81,260.0**														
Cotton	4,121.9	55,554.2	**531,125.7**													
Fonio	12,411.2	95,413.1	191,118.2	**672,753.8**												
Groundnut (High)	10,581.4	63,913.0	99,731.2	197,287.3	**393,740.4**											
Groundnut (Low)	9,338.7	57,207.1	100,503.3	201,462.0	102,038.5	**228,579.9**										
Maize (High)	-872.5	43,561.2	73,120.6	160,894.7	75,610.5	76,603.1	**265,225.9**									
Maize (Low)	3,520.1	22,916.6	70,770.8	85,695.3	66,882.7	59,507.5	66,895.1	**106,505.5**								
Okra	307.8	7,658.8	28,586.0	36,078.0	31,342.7	29,843.4	6,798.9	11,242.4	**22,530.8**							
Other	-3,884.0	11,394.5	173.0	70,563.1	15,908.7	33,983.6	21,794.3	7,009.8	5,272.0	**88,608.4**						
Pearl millet	6,232.4	40,035.8	107,265.6	145,214.9	85,505.9	84,187.9	55,230.0	50,501.1	13,101.2	18,333.2	**144,081.3**					
Rice	8,425.8	64,314.4	125,229.3	255,097.5	147,740.2	137,082.9	81,550.7	69,598.9	30,124.6	27,414.8	102,975.9	**378,348.0**				
Sesame	3,943.0	20,410.6	40,708.9	76,962.2	67,339.4	42,915.0	33,048.9	20,206.4	3,188.7	18,891.3	38,589.0	62,945.9	**105,944.6**			
Sorghum	5,604.8	38,134.2	79,635.6	121,135.1	78,066.0	70,630.8	59,244.2	44,352.3	11,863.8	19,543.9	62,467.2	78,329.4	31,520.8	**86,300.2**		
Sweet potato	5,601.3	29,968.6	46,114.8	78,815.9	56,398.3	46,114.6	17,962.3	9,809.3	18,364.8	21,746.3	22,361.7	44,417.8	39,190.9	21,502.1	**75,194.1**	
Yam	18,373.7	78,559.3	119,804.8	228,583.2	111,741.5	87,925.6	119,957.7	24,609.4	31,874.6	17,851.6	61,091.2	108,181.0	63,261.9	61,831.7	64,140.4	**322,719.4**

Notes: Lower triangular variance-covariance matrix of crop profits per hectare, reported in squared USD per hectare. Bold diagonal entries are the crop variances and run from 5,287.8 for Bambara groundnut to 672,753.8 for Fonio. Off-diagonal cells are shaded by the implied correlation rather than by the covariance itself. Blue indicates positive comovement and red negative, with intensity proportional to $|\rho|$. This is the matrix $\mathbf{C}$ that enters the portfolio variance $w'\mathbf{C}w$; Table B.7 reports the same object as correlations.

Table B.10: Realized and Predicted Profit Per Hectare Are Positively Correlated

	(1) No FE	(2) Year FE	(3) Crop-Set FE	(4) Crop-Set + Year FE
Expected Return/HA	1.563***	0.802***	1.539***	0.748***
	(0.059)	(0.052)	(0.059)	(0.052)
N	7,084	7,084	7,084	7,084
R^2	0.109	0.469	0.115	0.477
p (H_0: $\beta = 1$)	0.000	0.000	0.000	0.000
Crop-Set FE	No	No	Yes	Yes
Year FE	No	Yes	No	Yes

Notes: OLS regressions of realized profit per hectare on expected portfolio return. Columns vary by inclusion of year and crop-set fixed effects. A crop set is the menu of analysis crops cultivated in a community, so communities offering the same menu share one fixed effect and one efficient frontier; the sample contains 18 such sets. Positive coefficients indicate that portfolios with higher expected returns do achieve higher realized profits, validating the portfolio return measure. Standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Table B.11: Community Crop Availability: Share of Communities Missing Each Crop

Crop	**Communities Missing**	**% Missing**
Fonio	42	38.2%
Sesame	38	34.5%
Yam	34	30.9%
Bambara groundnut	5	4.5%
Pearl millet	5	4.5%
Cotton	2	1.8%
Sweet potato	2	1.8%
Beans	0	0.0%
Groundnut (high-intensity)	0	0.0%
Groundnut (low-intensity)	0	0.0%
Maize (high-intensity)	0	0.0%
Maize (low-intensity)	0	0.0%
Okra	0	0.0%
Other	0	0.0%
Rice	0	0.0%
Sorghum	0	0.0%

Notes: Table shows the number and percentage of communities (out of 110 total) where each crop is never observed being cultivated. 9 crops (Beans, Groundnut High, Groundnut Low, Maize High, Maize Low, Okra, Other, Rice, and Sorghum) are cultivated in all communities. Fonio, Sesame, and Yam are the most constrained, with 30.9–38.2% of communities lacking access.

Table B.12: Bivariate and Multivariate Associations Between Baseline Household Characteristics and Portfolio Efficiency

	Bivariate		Multivariate	
Predictor	$\hat{\beta}$ (SE)		$\hat{\beta}$ (SE)	
Uses Plough	0.174***	(0.016)	0.117***	(0.018)
Uses Fertilizer	0.094***	(0.015)	0.029*	(0.016)
Farm Size (HA)	0.127***	(0.017)	0.048**	(0.019)
Livestock Value (USD)	0.080***	(0.015)	0.018	(0.017)
Food Consumption per cap	0.002	(0.016)	0.004	(0.017)
Has Financial Savings	0.063***	(0.015)	0.038**	(0.015)
Household Size	0.140***	(0.015)	0.150***	(0.019)
Household Assets	0.021	(0.016)	0.071***	(0.019)
Agricultural Assets	-0.085***	(0.016)	-0.107***	(0.017)
Non-food Expenditures	0.100***	(0.017)	0.013	(0.019)
Not Enough Food for HH	-0.037**	(0.016)	-0.013	(0.016)
N	3,850–3,969		3,846	
Adj. R^2	—	—	0.115	—

Notes: Efficiency is proximity to the crop-set specific efficient frontier, signed so that higher values sit closer to the frontier. All variables are standardized, so each coefficient is in standard deviations of efficiency per standard deviation of the predictor. The bivariate columns report a separate regression of z(Efficiency) on each predictor alone, and are the estimates plotted in Figure A.16. The multivariate columns enter all eleven predictors jointly in a single regression (adjusted $R^2 = 0.115$). Both carry crop-set fixed effects, with standard errors clustered at household level, matching Figures A.16 and A.17. Both are estimated on the baseline (2010) survey round, before any transfers were made, so every covariate is pre-treatment and all study arms are pooled. Bivariate samples vary with item non-response, which is why N is reported as a range. These associations are correlational. Standard errors are heteroskedasticity-robust. * p<0.1, ** p<0.05, *** p<0.01.

Table B.13: Impacts of Cash Transfer on Measured and Estimated Outcomes

Panel A: Inputs, Output, and Scale

	Inputs & Output			Scale & Intensity		
	Total Expenditure	Total Revenue	Profit	Hectares Cultivated	Profit/HA	Expenditure /HA
Grant	30.97**	79.76*	74.44**	0.36***	4.49	0.59
	(14.57)	(43.57)	(32.31)	(0.11)	(3.16)	(1.55)
	[0.028]	[0.062]	[0.020]	[0.001]	[0.143]	[0.693]
N	6,276	6,276	6,276	6,276	6,276	6,276
Westfall-Young p	0.093	0.155	0.066	0.003	0.260	0.693
Control Mean	551.35	1625.40	1097.83	7.15	142.46	77.37
Treatment Mean	580.66	1672.48	1145.93	7.40	145.99	78.92

Panel B: Standard Proxies and MPT Outcomes

	Standard Proxies		MPT Outcomes		
	N Unique Crop Types	Normalized HHI	Expected Profit/HA	Portfolio St. Deviation	Sharpe Ratio
Grant	0.33***	-0.87**	-2.38	-4.01**	0.67**
	(0.07)	(0.43)	(1.76)	(1.83)	(0.26)
	[<0.001]	[0.044]	[0.144]	[0.016]	[0.009]
N	6,276	6,276	6,276	6,276	6,276
Westfall-Young p	<0.001	0.044	0.144	0.026	0.027
Control Mean	6.34	16.85	299.79	315.68	95.00
Treatment Mean	6.68	16.05	299.39	313.98	95.53

Notes: ANCOVA specification regressing each outcome on a treatment indicator (grant receipt), the baseline value of the outcome, the authors' preferred control suite, and a year fixed effect. Standard errors clustered at household level in parentheses. Randomization inference p-values (10,000 permutations) in brackets. Sample restricted to grant-only villages (arm 4) versus pure control villages (arm 5), follow-up rounds 1–2 (2011–2012), balanced panel households only. All monetary values in 2017 USD. * p<0.1, ** p<0.05, *** p<0.01. The Westfall-Young row reports stepdown adjusted p-values controlling the familywise error rate within outcome families rather than across the table (see Appendix E). This table spans three of the four families: hectares, crop count and concentration form the ex-ante family; expenditure, revenue, profit, profit per hectare and expenditure per hectare the ex-post family; expected profit per hectare and portfolio standard deviation the MPT family. The fourth is adjusted in Table B.14.

Table B.14: Impacts of Cash Transfer on Portfolio Efficiency

	Sharpe Ratio	Euclidean Dist, Crop-Set EF	Euclidean Dist, Crop-Set k-EF	Sharpe at Ideal Point	Selection Ratio	Diversification Ratio
Grant	0.67**	-1.48***	-0.25	0.09	0.68	1.16***
	(0.26)	(0.53)	(0.47)	(0.11)	(0.95)	(0.35)
	[0.009]	[0.005]	[0.579]	[0.389]	[0.435]	[<0.001]
N	6,276	6,276	6,276	6,276	6,276	6,276
Westfall-Young p	0.027	0.017	0.739	0.739	0.739	0.004
Control Mean	95.00	49.29	39.38	119.63	179.56	135.43
Treatment Mean	95.53	48.11	39.43	119.62	178.91	136.47

Notes: Same specification as Table B.13. Sharpe Ratio is the ratio of expected return to portfolio standard deviation, scaled by 100. Euclidean Dist to Crop-Set EF is the distance to the closest point on the crop-set specific efficient frontier; the Crop-Set k-EF variant restricts to the cardinality-k frontier with k equal to the household's number of cultivated crops. Sharpe at Ideal Point is the Sharpe ratio of the closest efficient frontier portfolio itself. It is a level, not a difference from the household's own Sharpe ratio, which is why its control mean (119.6) exceeds the control mean of the household's own Sharpe ratio (95.0). Selection Ratio is the standard deviation of the community's highest-return crop divided by the household's linear-weighted average crop standard deviation, $\sum_i w_i \sigma_i$ (Equation 6). Diversification Ratio compares actual portfolio SD to the no-diversification benchmark (Equation 5). Standard errors clustered at household level in parentheses. Randomization inference p-values (10,000 permutations) in brackets. The Westfall-Young row reports stepdown adjusted p-values controlling the familywise error rate within the efficiency family (see Appendix E), which contains every column shown here and one further outcome, origin distance [0.146], reported in Appendix H. Because the stepdown enforces monotonicity in the sorted order, the least significant members of a family share a common ceiling, which is why three columns here carry the same adjusted value. * p<0.1, ** p<0.05, *** p<0.01.

Table B.15: Heterogeneous Treatment Effects by Baseline Wealth Tercile, with Pairwise Contrasts

	Treatment effect within tercile			Pairwise contrast (SE)		
Outcome	T1 (poorest)	T2 (middle)	T3 (richest)	T1 − T2	T3 − T2	T3 − T1
Number of crops	+0.1	+0.4***	+0.4***	−0.3* (0.2)	+0.0 (0.2)	+0.3* (0.2)
Expected return / HA	−3.7	+0.3	−3.0	−3.6 (4.3)	−2.9 (4.2)	+0.7 (4.3)
Portfolio SD / HA	−6.0*	+0.2	−5.1	−6.1 (4.5)	−5.3 (4.5)	+0.8 (4.5)
Sharpe ratio (×100)	+1.0**	+0.1	+0.9**	+0.9 (0.7)	+0.9 (0.6)	−0.1 (0.6)
Distance to community frontier	−2.3**	−0.3	−1.7**	−2.1 (1.4)	−1.5 (1.3)	+0.6 (1.3)
Distance to community k-frontier	−1.8**	+1.4*	−0.3	−3.3*** (1.2)	−1.5 (1.1)	+1.8 (1.1)
Selection Ratio (self-insurance)	+2.8	−2.2	+1.5	+4.8** (2.3)	+3.4 (2.2)	−1.3 (2.4)
Diversification Ratio	+0.7	+1.4**	+1.2**	−0.8 (0.9)	−0.2 (0.8)	+0.6 (0.8)
n (HH-rounds)	2,084	2,084	2,085		6,253	

Notes: The left panel reports the grant coefficient from the headline ANCOVA estimated separately within baseline wealth-index terciles (T1 poorest, T3 richest). The right panel reports pairwise differences between those effects, estimated from a pooled model interacting the grant indicator with tercile, with the standard error of the difference in parentheses. The index is an unweighted sum of nine z-scored baseline endowment variables (Section 4.4). Controls and SE clustering as in Table B.13. Randomization was stratified on village rather than on baseline wealth, so the experiment was not designed to power comparisons across terciles: the minimum detectable difference in portfolio standard deviation between terciles is roughly twice the size of the average effect itself. The contrasts should therefore be read as describing the pattern of responses rather than as tests the design was built to pass. Two coefficients discussed in the text as suggestive fall just outside conventional significance: the T1 Selection Ratio increase ($p = 0.113$) and the T3 portfolio standard deviation reduction ($p = 0.105$). * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Appendix C: The Profit Measure: Construction and Robustness

The foundation of this work is the high resolution survey data that enables plot-crop level measures of profits per hectare. Without accurate profit measures it would not be possible to estimate the variance-covariance matrix of returns. There are multiple choices made while calculating the profit measures and they are answers to the following questions

1. What elements of inputs should be included in the calculation?
2. How to value the inputs included in the calculation?
3. How to handle costs on intercropped plots?
4. How to value the output of a plot?

Each one of these decisions is consequential, and this appendix details the process and researcher choices; Table C.1 in Section C.1 reports the main-suite effects under the alternative convention, an equal split of an intercropped plot's area and costs across its crops.

C.1: Cost and Area Allocation Conventions for Intercropped Plots

Cost allocation. The main specification allocates plot-level input costs to crops on intercropped plots in proportion to each crop's share of plot revenue. This assumes that higher-value crops receive proportionally more inputs. As an alternative, the full pipeline can be re-estimated under an equal-split cost allocation, where each crop on an intercropped plot receives $1/N$ of the plot's total costs regardless of revenue. Table C.1 compares the two conventions across the 7 main-suite outcomes, each rebuilt from the plot level under the equal split: crop moments, the covariance matrix, portfolio weights and outcomes all change, so the comparison is of two complete constructions rather than of one input. Every effect keeps its sign, and the area, crop-count and profit effects are essentially unchanged. The portfolio-SD effect is larger under the equal split, -8.49 against -4.01 USD/HA, with a standard error about twice the headline's (3.90 against 1.83); part of the gap is scale, since the equal split gives low-revenue crops more area and raises the control-mean portfolio SD from 316 to 423 USD/HA, but the effect is larger relative to its control mean as well (2.0% against 1.3%). The revenue-share convention is therefore the conservative one for the headline result.

Area allocation. The same revenue-share convention used for costs also governs how plot area is allocated across constituent crops on an intercropped plot, so that the per-crop hectare weights w_c sum to one within a household-year as the MPT framework requires. An alternative convention used in some crop-adoption studies counts the full plot area toward each crop separately, which can produce $\sum_c w_c > 1$ but answers a different question (*how much land has crop c on it?* rather than *how much land is attributable*

Table C.1: Robustness: Revenue-Share vs. Equal-Split Allocation of Intercropped Plots

	Revenue share (headline)		Equal split	
	Coef	SE	Coef	SE
Hectares	0.36***	(0.11)	0.31***	(0.10)
N Crops	0.33***	(0.07)	0.34***	(0.07)
Norm HHI	-0.87**	(0.43)	-0.81*	(0.42)
Profit	74.44**	(32.31)	77.60**	(34.70)
Profit/HA	4.49	(3.16)	4.18	(4.66)
E[Profit]/HA	-2.38	(1.76)	-6.47*	(3.82)
Portfolio SD	-4.01**	(1.83)	-8.49**	(3.90)

Notes: Grant treatment effects under the headline convention, which allocates an intercropped plot's area and input costs to its crops by revenue share, and under an equal split of both across the crops on the plot. The equal-split column is a full rerun of the construction pipeline under that convention: crop moments, the covariance matrix (with the same $N - G$ degrees-of-freedom correction), portfolio weights and outcomes are all rebuilt, so every outcome can differ, not only the profit measures. Both columns use the headline ANCOVA specification: the authors' control suite, round indicators, a baseline control for the dependent variable, and standard errors clustered at the household. * p<0.10, ** p<0.05, *** p<0.01.

to crop c?). Empirically, the alternative convention overstates crop-specific area for approximately 9% of household-years in this sample on the rice + groundnut margin, but it cannot be used here because it breaks the portfolio-weight identity required for MPT optimization. The revenue-share allocation also has a known potential bias: when a crop fails on an intercropped plot, the allocator $\text{rev}_c / \sum_j \text{rev}_j$ co-varies with the realized state of the world. The most direct robustness for the area-allocation question is reported in Appendix D.3, which re-estimates the variance-covariance matrix $\mathbf{C}$ at the community-year level (where farmer-specific allocation idiosyncrasies wash out in the cross-farmer average); the headline portfolio-SD treatment effect is essentially identical under the community-level estimation ($-\$3.62$, $p = 0.020$) as under the farmer-level estimate from the same specification ($-\$4.01$, $p = 0.029$).

Within-plot mechanical correlation in profit per hectare. A related concern is that, on intercropped plots, allocating both costs and area by revenue share yields algebraically identical profit per hectare across all crops on the plot. With $A_c = s_c A$ and $C_c = s_c C$ where $s_c = R_c/R$, the per-hectare profit for crop c on the plot reduces to $(R - C)/A$, independent of which crop is being valued. Within-plot profit-per-hectare residuals across crops are therefore mechanically perfectly correlated, and pairs of crops that are frequently intercropped together would have inflated correlation estimates if all of the variance-covariance evidence came from co-located observations. Two empirical features attenuate this concern. First, of the 36.2% of plot-crop observations that are sole-cropped, allocation is unnecessary and profit-per-hectare differentiation across crops is uncontaminated by the mechanical equality. Second, when crop-pair correlations are aggregated to the farmer-year level (the unit of estimation for $\mathbf{C}$), only 22.1% of farmer-year crop-pair observations involve crops co-located on at least one shared plot; the remaining 77.9% of pair observations come from cross-plot comparisons within a farmer-year, where each crop's profit per hectare is informed by

separate plot-level realizations. Appendix D.4 goes further than this prevalence argument and re-estimates every correlation on cross-plot pairs alone, leaving the headline risk effect intact.

C.2: Sensitivity to Crop Classification (Dip Test Splits)

The main specification uses Hartigan's Dip Test to identify crops with multi-modal cost distributions and splits them into high-intensity and low-intensity variants, treating each as a separate asset class (Section 2.2). Table C.2 tests the sensitivity of crop-level treatment effects to this classification by comparing the grant's effect on the high-intensity share, low-intensity share, and merged (unsplit) share for each affected crop. If the merged coefficient is similar to the weighted combination of the split coefficients, the dip test classification does not materially affect the results.

Table C.2: Dip Test Sensitivity: Treatment Effects on Split vs. Merged Crop Shares

Crop	Variant	Coefficient	SE	p-value
Groundnut	High intensity	0.0048**	(0.0023)	0.033
	Low intensity	-0.0058**	(0.0027)	0.033
	Merged	-0.0009	(0.0033)	0.772
Maize	High intensity	-0.0009	(0.0028)	0.749
	Low intensity	-0.0113	(0.0071)	0.110
	Merged	-0.0122*	(0.0072)	0.091

Notes: Treatment effects of the grant on crop acre shares, comparing high-intensity and low-intensity variants (from Hartigan's Dip Test classification) with the merged share (no split). Same ANCOVA specification as main results. Standard errors clustered at household level. * p<0.10, ** p<0.05, *** p<0.01.

C.3: Measurement Error Across Crops

A potential concern with the MPT framework is that measurement error in profit per hectare could be systematically different across crops. If so, a treatment that shifts crop composition could change the measurement error component of portfolio variance, confounding the treatment effect on true risk. As discussed in Section 2.3, the threat requires measurement error to be systematically larger for some crops than others, so that shifting into those crops would inflate measured variance even without a change in true risk.

To test this directly, I adapt the within-farmer decomposition strategy of Gollin and Udry (2021). Many farmers in Sikasso grow the same crop on multiple plots within a single growing season. Under the assumption that within-farm allocation is efficient, any variation in profit per hectare across plots of the same crop, grown by the same farmer, in the same season, can be attributed to three sources: (1) measurement error in inputs and outputs, (2) late-season production shocks realized after input decisions, and (3) unobserved land quality

heterogeneity across plots. The within-farmer, within-crop, within-season variance of profit per hectare is therefore an upper bound on the measurement error variance.

In the pure-control sample that estimates $\mathbf{C}$ (2,390 farmers in 110 communities), 49% of farmers grow at least one crop on two or more sole-cropped plots within a season, yielding 2,759 farmer-crop-season cells. The bound is taken on sole-cropped plots only. On an intercropped plot the revenue-share allocation of Appendix C.1 gives every constituent crop the plot's profit per hectare, so two plots of the same crop on the same farm can differ because one is sole-cropped and the other shares a plot with another crop; that is composition dispersion, not measurement error, and Table C.3 reports it separately as a contrast. For each crop, I compute the ratio of the mean within-cell variance to the total crop variance from the covariance matrix diagonal. This ratio, the "measurement error share", captures the fraction of total crop variance attributable to within-farmer variation (an upper bound on pure measurement error).

Before reporting the shares, one feature of the variance estimator deserves comment, because it bears directly on whether the shares are comparable across crops. The deviations entering each crop's variance are taken around farmer-crop means that are themselves estimated from the same observations, so within a cell of T_g seasons the deviations sum to zero and carry only $T_g - 1$ independent pieces of information. Summing across the G_c cells in which crop c is grown, the expected sum of squared deviations is $(N_c - G_c)\sigma_c^2$ rather than $N_c\sigma_c^2$, and the residual degrees of freedom are therefore $N_c - G_c$. Using the conventional $N_c - 1$ divisor in its place understates each crop's variance, and it does so unevenly: the attenuation is approximately $1 - 1/\bar{T}_c$, where $\bar{T}_c$ is the average number of seasons per farmer-crop cell. In this panel $\bar{T}_c$ ranges from 2.3 for sweet potato to 5.6 for rice, so the corrected variance is larger than the conventional one by roughly 22% to 76% depending on the crop. That is not a common scale factor, and a common scale factor is what a cross-crop comparison would need in order to be indifferent to the choice. The variances reported throughout the paper accordingly use the $N_c - G_c$ divisor. The correction leaves the estimated correlation matrix untouched, since correlations are invariant to any common rescaling of the two series entering them, so it affects the magnitude of crop risk without disturbing the comovement structure on which the portfolio results depend.

Table C.3 reports the results. The estimated upper bounds on the measurement error share range from 13% to 46% among the 6 major crops, with a mean of 22% and a coefficient of variation of 0.54. These bounds stay under half for every major crop; they are not uniform, and uniformity is in any case not what identification requires. On intercropped plots the same ratio runs from 39% to 69% (mean 50%), and the excess over the sole-cropped bound is the plot-composition dispersion that the allocation convention writes into the diagonal of $\mathbf{C}$: a footprint of Appendix C.1, not of the survey instrument. Because treatment is randomized and the survey instrument is identical across arms, measurement error is balanced between

treatment and control within every crop, so level bias is also balanced across arms. Two features of the design then bound what the remaining channel, treatment-induced composition shifts across crops with different error loads, can contaminate. First, the cross-crop ordering of risk survives the correction that differential mismeasurement requires. Writing the observed variance as signal plus noise and letting $s_c = 1 - m_c$ denote the signal share of crop c, the observed ranking of a pair reproduces the true ranking exactly when $\hat{\sigma}^2_A/\hat{\sigma}^2_B > s_B/s_A$: the threshold sits near one for crops carrying similar error loads and rises only where the apparently riskier crop is also the noisier one. Applying this test to every pair of major crops, 14 of 15 orderings survive (93%), and the most demanding threshold any pair must clear is 1.12. The one ordering that does not is groundnut (high-intensity) over rice (observed variance ratio 1.04 against a threshold of 1.09). Extending the test to the full crop set, 89% of orderings survive; of the 13 that do not, 12 involve a crop below the cells floor of the table, where the within-cell variance is too thinly estimated to bound anything. Because the within-farmer variance is an upper bound on measurement error, the signal shares used here are lower bounds and the test understates how many orderings survive.

Second, plot-level measurement error is independent across plots and crops, so it inflates the diagonal of the covariance matrix but cannot generate off-diagonal comovement; the hedging channel that Section 6 identifies as carrying the treatment effect operates through precisely those off-diagonal terms, which this class of measurement error cannot manufacture. The within-farmer variance used here remains an upper bound on measurement error, since it also contains genuine late-season shocks and land-quality differences; a bound tight enough to be quantitative would require separating those components, which these data cannot do.

Table C.3: Measurement Error Share by Crop

	Sole-cropped plots (the bound)				Intercropped plots	
Crop	Cells	Total Var	Within Var	ME Share	Cells	Share
Rice	1,061	378,348	60,986	16%	1,285	44%
Groundnut (low-intensity)	446	228,580	37,828	17%	2,024	44%
Groundnut (high-intensity)	196	393,740	89,421	23%	406	69%
Sorghum	170	86,300	15,269	18%	637	46%
Maize (low-intensity)	148	106,506	14,350	13%	232	39%
Beans	125	81,260	37,026	46%	775	55%
Sole-cropped bound across all crops: Mean = 22%, Median = 16%, CV = 1.09						

Notes: "Cells" are farmer-crop-season observations with 2+ sole-cropped plots of the same crop. "Total Var" is the diagonal of the covariance matrix (estimated from the control group). "Within Var" is the average within-cell variance of profit/HA across those plots. The intercropped columns repeat the calculation on intercropped plots and measure the plot-composition dispersion of the revenue-share allocation (Appendix C.1), not measurement error. "ME Share" = Within Var / Total Var. High- and low-intensity variants are separate crops here, as in the covariance matrix. Table restricted to the 6 major crops, those with at least 500 baseline cultivators and at least 100 sole-cropped multi-plot cells in the pure-control sample; the "other" aggregate is excluded. All values in 2017 USD. Method adapted from Gollin and Udry (2021).

C.4: Measurement Error at the Portfolio Level

Appendix C.3 asks whether measurement error is uniform enough across crops to preserve their risk ranking. That argument is about the ranking of individual crops, and it does not carry over to the treatment effect on portfolio-level variance, which is the paper's headline outcome. Under uniform additive measurement error with variance V_{ME} per crop independent across crops, observed portfolio variance is true portfolio variance plus a concentration-weighted measurement-error term: $\widehat{\sigma}_p^2 = \sigma_p^2 + V_{\text{ME}} \sum_i w_i^2$, where $\sum_i w_i^2$ is the Herfindahl-Hirschman index (HHI) of portfolio concentration. The treatment-control difference therefore includes a mechanical $V_{\text{ME}} \cdot \Delta\text{HHI}$ term in addition to the true-variance treatment effect. Because the grant reduces concentration ($\Delta\text{HHI} < 0$), this term is negative, and the estimated variance reduction is biased toward larger magnitudes. The magnitude of this mechanical term cannot be pinned down credibly: the within-farmer variance that would parameterize V_{ME} is an upper bound contaminated by genuine late-season shocks, and Appendix C.3 details why these data cannot separate the two. Its direction, however, is known, so the headline variance effect reported in Section 6 should be read as an upper bound on the true-variance reduction. The off-diagonal hedging component of the effect, which Section 6 identifies as the operative channel, is immune to this term entirely, because plot-level error cannot generate cross-crop comovement. This bias is one of the reasons the paper reports portfolio variance both in the squared-weight diagonal form ($\sum_i w_i^2 \sigma_i^2$, which is exposed to it) and in the linear-weighted selection ratio of Equation 6, which does not carry the same $V_{\text{ME}} \cdot \text{HHI}$ contamination.

C.5: Revenue-Based vs Profit-Based Portfolio SD

The headline portfolio-SD result uses crop-level moments computed from per-hectare profits (revenue net of allocated input costs), as detailed in Sections 2.2 and 2.3. A natural robustness question, raised informally by audience members on multiple occasions, is whether the choice of dependent variable, profits versus revenues, drives the finding. To answer this, I rebuild the covariance matrix from per-hectare revenues using the same farmer-demeaning + pure-control pipeline as the main specification, apply the existing portfolio weights to construct a revenue-based portfolio standard deviation, and re-estimate the headline treatment effect.

Table C.4: Treatment effect on portfolio SD: profit-based vs revenue-based

	Profit-based PortSD ($/HA)	Revenue-based PortSD ($/HA)
Grant	-4.011**	-3.175
	(1.832)	(1.941)
Control mean	315.68	309.20
TE as % of ctrl mean	-1.27%	-1.03%
N	6,276	6,276

Notes: Both columns use the headline ANCOVA specification of Section 6: the Beaman et al. (2023) controls, round indicators, a baseline control for the dependent variable, and standard errors clustered at the household. The profit-based column reproduces the headline result from Table B.13. The revenue-based column substitutes a covariance matrix built from per-hectare revenues for the matrix built from profits, holding portfolio weights and treatment-effect specification fixed. * p<0.10, ** p<0.05, *** p<0.01.

The revenue-based effect (Table C.4) is similar in magnitude to the profit-based effect but less precisely estimated: -3.17 USD per hectare ($p = 0.102$, control mean \$309) versus -4.01 ($p = 0.029$, control mean \$316). Expressed as a percentage of each control mean, the two effects are similar (1.0% and 1.3%). The substantive reading is that the sign and size of the headline portfolio-SD result do not depend on the choice of dependent variable, while its precision does: revenues carry more measurement noise than profits net of allocated costs, and the standard error rises accordingly. Both representations load on the household's exposure to the chosen crop mix's intrinsic variance, which is largely shared between revenue-only and profit-based measures of crop returns. The choice of profits as the headline dependent variable, motivated in Section 2.2 by the cost heterogeneity that revenues alone cannot capture, does not appear to be load-bearing for the headline finding.

C.6: Bambara Groundnut and the Baseline Roster Omission

One crop enters the panel on a different footing from the rest, and because it sits at the minimum-variance corner of most community frontiers it deserves its own test. The 2010 baseline questionnaire's crop roster runs from 1 (rice) to 27 (other) and contains no code for Bambara groundnut. The follow-up and endline rosters were renumbered and do carry the crop, as code 26, *pois de terre*. Bambara groundnut therefore has a mechanically zero baseline share for every household, and its σ_c is estimated without the 2010 season while every other crop's is estimated with it. Since dropping a season removes cross-sectional variation, this understates its variance relative to the rest of the crop set, and an understated variance is precisely what pulls a crop toward the minimum-variance corner of the frontier.

The natural robustness exercise is not to delete the crop. Deleting it would drop those household-years'

area and profit from the estimation sample as well, so the resulting movement could not be attributed to the taxonomy alone. Instead I apply the baseline roster convention to every round: *pois de terre* is folded into the residual "other" category, exactly as the 2010 instrument would have recorded it. The fold preserves every household's cultivated area and profit, so no observation leaves the estimation sample. What it does change is the taxonomy: the menu falls from 16 assets to 15, and the residual category is re-estimated with the folded plots inside it, which lowers its standard deviation from \$298 to \$284 per hectare. That is the counterfactual the 2010 instrument implies rather than a side effect of it, and it is why the frontier moves below while the treatment effect does not.

Table C.5 reports both arms. The fold does move the frontier, as the diagnostic above predicts. Control-group distance to the crop-set efficient frontier falls from 49.29 to 47.93, a shift of 1.36, because the crop that was anchoring the minimum-variance corner is no longer a separate asset. What does not move is the treatment effect. The grant's effect on portfolio standard deviation, the paper's headline outcome, goes from −\$4.01 per hectare (SE 1.83) to −\$3.98 (SE 1.82). The grant's effect on distance to the frontier goes from −1.48 (SE 0.53) to −1.44 (SE 0.51), a change of 2.7%, and remains significant at the one percent level. The Sharpe ratio and diversification ratio effects are likewise unchanged to within a hundredth (0.67 versus 0.66, and 1.16 versus 1.17).

That the level moves while the effect does not is the substance of the test rather than an incidental detail. A crop with a damaged baseline can distort where the frontier sits, and here it demonstrably does. For it to threaten the paper's conclusion it would have to distort how far the grant moves households relative to that frontier, and it does not. The headline result is a statement about the treated-control difference, and that difference is invariant to whether Bambara groundnut is treated as its own asset or absorbed into the residual category.

C.7: Does the Risk Measure Order Realized Variation?

Table B.10 validates the expected-return construct by regressing each household's realized profit per hectare on the return its crop mix predicts. No comparable exercise validates $\sigma_p = \sqrt{\vec{w}'\mathbf{C}\vec{w}}$, which every result in the paper runs through. This section supplies one, at a weaker standard than the return exercise meets and deliberately so.

The obstacle is that a single season's outcome says nothing about dispersion. Realized return can be compared to predicted return household by household, because each household supplies one of each. Risk cannot: a household's realized variation is only visible across seasons, and four seasons give three departures from its own average, which is far too few to estimate a household-specific volatility. The comparison

Table C.5: Treatment Effects With and Without Bambara Groundnut Folded into "Other"

	Expected Profit/HA	Portfolio St. Deviation	Sharpe Ratio	Euclidean Dist, Crop-Set EF	Euclidean Dist, Crop-Set k-EF	Sharpe at Ideal Point	Selection Ratio	Diversification Ratio
No fold (reproduction)								
Grant	-2.38	-4.01**	0.67**	-1.48***	-0.25	0.09	0.68	1.16***
	(1.76)	(1.83)	(0.26)	(0.53)	(0.47)	(0.11)	(0.95)	(0.35)
Control mean	299.79	315.68	95.00	49.29	39.38	119.63	179.56	135.43
Bambara groundnut folded into "other"								
Grant	-2.39	-3.98**	0.66***	-1.44***	-0.19	0.12	0.66	1.17***
	(1.75)	(1.82)	(0.25)	(0.51)	(0.45)	(0.11)	(0.94)	(0.34)
Control mean	300.22	316.09	94.99	47.93	38.43	118.62	179.37	135.42
Difference	-0.01	0.03	-0.01	0.04	0.06	0.03	-0.02	0.01
N	6,276	6,276	6,276	6,276	6,276	6,276	6,276	6,276

Notes: Each arm is a complete run of the estimation pipeline from the raw survey data, identical in every respect except the treatment of *pois de terre*. In the upper panel the crop is its own asset, as in the main text; in the lower panel it is folded into the residual "other" category, applying the 2010 baseline roster convention to all rounds. The first two columns are the headline MPT outcomes of Table B.13 and the remaining six the efficiency outcomes of Table B.14; the upper panel reproduces both exactly. Specification, controls, fixed effects and sample match the main efficiency suite. Standard errors, clustered at the household level, in parentheses. $N = 6,276$ in both arms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

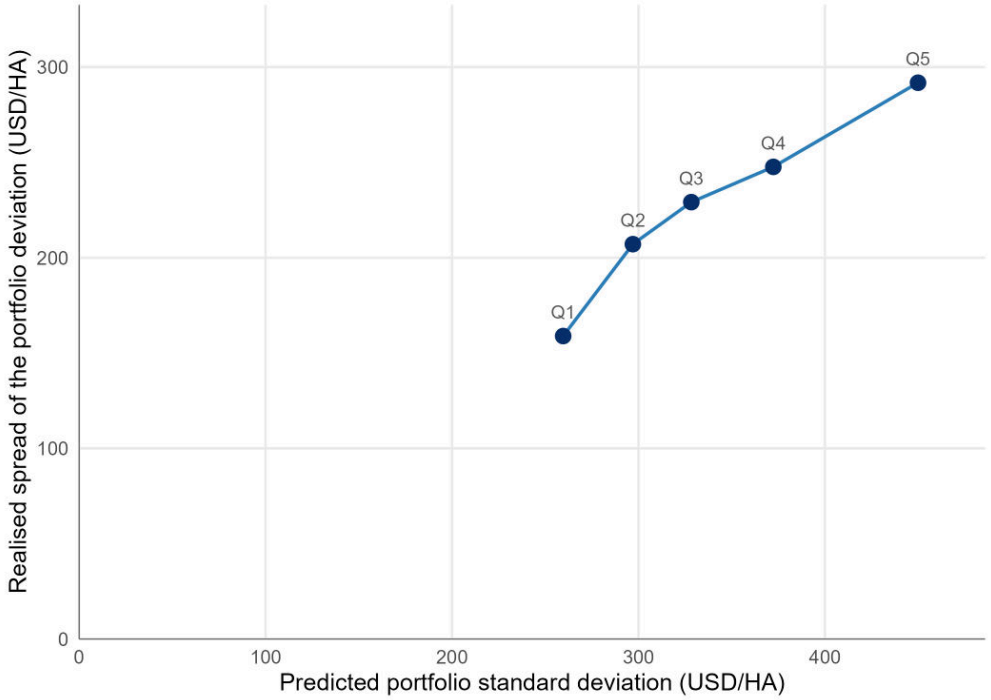
therefore has to be made across groups of households rather than within them, and what it can establish is an ordering rather than a magnitude.

The exercise is four steps. For each crop a household cultivated, I take the gap between its profit per hectare that season and the household's own average profit per hectare in that crop, which is the deviation from which $\mathbf{C}$ is estimated. I combine those gaps using the household's area shares, giving a portfolio deviation in dollars per hectare for each household-season. I then sort household-seasons into quintiles of predicted σ_p and compare the spread of the portfolio deviation across quintiles. No adjustment is applied at any stage. The sample is the 9,302 pure-control household-seasons of the 2,375 farmers whose realizations estimate the moments; the treated arm is not used, since its post-grant seasons would confound a measurement question with the treatment itself.

Figure C.1 and Table C.6 report the result. Predicted risk rises 1.73-fold from the bottom quintile to the top, from \$260 to \$450 per hectare, and the realized spread rises 1.84-fold alongside it, from \$159 to \$292. The ordering is strict across all five quintiles, and across ten bins the rank correlation between predicted and realized is 0.98. Households the framework assigns more risk did bear more realized variation.

Two limits are worth stating plainly. The exercise establishes ordering and not calibration: it does not show that a household predicted to face \$400 of risk per hectare faced \$400 rather than \$300. And the households scored here are the same ones whose realizations estimate $\mathbf{C}$, so they are not a held-out sample; with 2,375 farmers contributing to each moment the resulting circularity is of order $1/N$, but it is not zero. Appendix C.9 removes the same overlap for the return construct and finds it immaterial. What the section rules out is the possibility that the risk measure ranks households arbitrarily, which is the property every treatment effect in the paper relies on, since each compares two arms priced by the same matrix.

Figure C.1: Realized Variation Rises with Predicted Portfolio Risk



Notes: Pure-control household-seasons sorted into quintiles of predicted portfolio standard deviation, labelled Q1 to Q5. The vertical axis is the standard deviation, within the quintile, of the household-season portfolio deviation defined in the text. Both axes begin at zero, so the two rises are shown in proportion. The axes are not calibrated against one another and no forty-five degree reference line is drawn, because only the ordering is interpretable: no degrees-of-freedom or scaling adjustment has been applied to either series. Predicted risk in this sample does not fall below roughly $260 per hectare, which is why the points occupy the right of the horizontal range.

Table C.6: Predicted Risk and Realized Variation, by Quintile

Quintile of predicted risk	Predicted SD	Realised spread	Household-years
1	260	159	1,861
2	297	207	1,861
3	328	229	1,860
4	372	248	1,860
5	450	292	1,860

Notes: Pure-control household-seasons sorted into quintiles of predicted portfolio standard deviation. "Predicted SD" is the mean of σ_p within the quintile. "Realized spread" is the standard deviation, within the quintile, of the household-season portfolio deviation: the household's area-weighted sum of the gaps between each crop's realized profit per hectare and that household's own average for the crop. Both columns are in USD per hectare. No degrees-of-freedom or scaling adjustment is applied to either column, so the two are not calibrated against one another and only their ordering is interpretable.

C.8: Do Control Moments Price Treated Portfolios?

Every household in the paper, treated or not, is priced at the moments estimated from pure-control farmers.

That is innocuous only if grant recipients realize, from the same crops, what those moments predict for their

weights. If the grant led households to farm a given crop more intensively, its true return and risk for them would differ from $\bar{R}_c$ and σ_c, and the portfolio measures, which move only when weights move, would not see it. Expenditure per hectare did not move (Section 6), which bounds the channel on the input side; this section tests it on the outcome side, at the level of the household portfolio, so that no crop-level moment has to be re-estimated on the 797 treated farmers.

The return test extends the validation regression of Table B.10 to both arms in the two post-treatment rounds, adding an indicator for grant receipt and its interaction with the expected return, which is centered at its control mean so that the indicator is the level shift at the mean portfolio. Table C.7 reports it. Treated households realize \$3.75 per hectare more than their weights predict (SE 3.34, $p = 0.262$), a 95 percent interval of -0.9 to 3.4 percent of the \$300 mean expected return; the slope is common to both arms (interaction 0.01, $p = 0.807$), and the two grant terms are jointly zero ($p = 0.528$). Whatever treated households do differently with a given crop is worth at most a few percent of the return the framework assigns them.

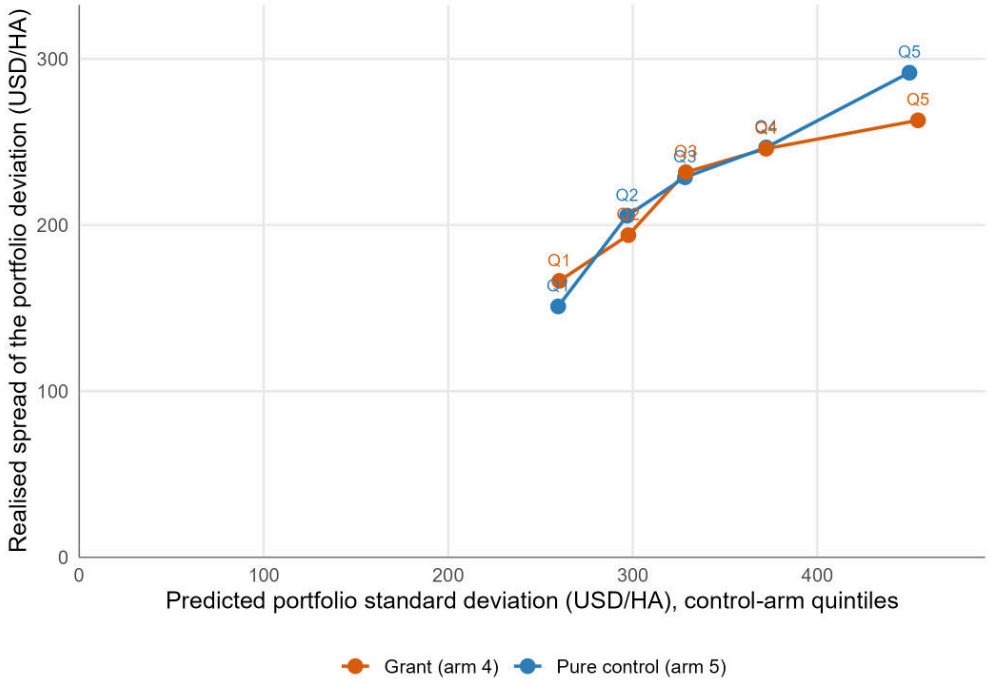
The risk test repeats the ordering exercise of Appendix C.7 on both arms, with the deviations of both built by one rule and the household-seasons binned on the control arm's quintile cut points, so that the arms are compared at matched predicted risk. Figure C.2 and Table C.8 report it. Treated household-seasons sit on the control curve: realized spread rises 1.58-fold across the quintiles against 1.93-fold for control, the ordering is strict, and the decile rank correlation is 0.87 against 0.98 for control. Regressing the absolute portfolio deviation on a grant indicator with quintile and year effects, treated households differ from control by -2.0 dollars per hectare of realized spread at the same predicted risk (SE 4.3, $p = 0.644$; -1.2 percent of the control mean), the arm-by-quintile interactions are jointly zero ($p = 0.252$), and restricting both arms to post-grant seasons gives 0.7 dollars (SE 4.1, $p = 0.864$). Control-estimated moments price treated portfolios in return and in risk, which is what every treatment effect in Section 6 assumes when it compares two arms priced by one matrix.

Table C.7: Realized Profit per Hectare against Predicted Return, Both Arms

	(1) Year FE	(2) Crop-Set + Year FE
Expected return/HA (centred)	0.513***	0.513***
	(0.031)	(0.031)
Grant recipient	3.75	3.75
	(3.34)	(3.34)
Grant × expected return/HA	0.015	0.015
	(0.061)	(0.061)
p, grant terms jointly zero	0.528	0.528
N	6,282	6,282
R^2	0.113	0.113
Crop-Set FE	No	Yes
Year FE	Yes	Yes

Notes: OLS of realized profit per hectare on the framework's expected return (control-estimated crop moments applied to each household's own land shares, centered at its control-arm mean of $300), an indicator for grant receipt, and their interaction. Grant (arm 4) and pure-control (arm 5) households in the two post-treatment rounds: 6,282 household-years, 1,579 of them treated. The grant coefficient is the dollars per hectare that treated households realize above what their weights predict at the mean portfolio; the interaction tests whether the cross-crop structure differs by arm. The slope is below that of Table B.10 because this sample excludes the baseline round; the interaction shows it is common to both arms. Standard errors clustered at the household level in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Figure C.2: Realized Variation against Predicted Risk, Both Arms



Notes: Household-seasons of both arms sorted into quintiles of predicted portfolio standard deviation using the control arm's cut points, so the two curves are compared at matched predicted risk. The vertical axis is the within-quintile standard deviation of the household-season portfolio deviation, built for both arms by one rule: each farmer-crop mean is taken over every plot-crop observation of that farmer in that crop, cells need at least two seasons, and deviations are averaged within a farmer-season-crop before being combined with land shares. The control curve reproduces Appendix C.7 up to that appendix's rice-matching filter. 9,300 control and 3,121 treated household-seasons; both axes start at zero.

Table C.8: Predicted Risk and Realized Variation by Quintile, Both Arms

	Pure control (arm 5)			Grant (arm 4)		
Quintile of predicted risk	Predicted SD	Realised spread	HH-years	Predicted SD	Realised spread	HH-years
1	260	151	1,860	260	166	659
2	297	206	1,860	298	194	601
3	328	229	1,860	329	232	609
4	372	247	1,860	372	246	593
5	450	292	1,860	455	263	659

Notes: Quintiles are cut on the control arm's predicted portfolio standard deviation and applied to both arms. "Predicted SD" is the mean of σ_p within the quintile and arm; "Realised spread" the standard deviation of the household-season portfolio deviation within it. Both columns are in USD per hectare; as in Table C.6, no degrees-of-freedom or scaling adjustment is applied, so only the ordering and the comparison across arms at a given quintile are interpretable.

C.9: Leave-One-Out Validation of the Return Construct

Table B.10 scores each control household-year's realized profit per hectare against $\vec{w}'\bar{R}$, and every control household that grew a crop in at least two seasons contributes its own farmer-crop mean to that crop's

$\bar{R}_c$. The regressor is therefore not independent of the outcome. With 130 to 2,049 farmer-crop cells behind each crop mean the overlap is of order $1/N$, but it need not be conceded, and this section removes it. For each control household h and crop c the expected return is rebuilt from a crop mean that excludes h's own farmer-crop mean, $\bar{R}_c^{(-h)} = (N_c\bar{R}_c - m_{hc})/(N_c - 1)$; a household that never contributed to crop c keeps $\bar{R}_c$; and the four specifications of Table B.10 are re-estimated on the identical 7,084 household-years with the leave-one-out predictor. 99 percent of household-years contributed to at least one crop mean, so almost every predictor changes, but it changes by 0.09 dollars per hectare on average (correlation 1.000 with the full-sample predictor). Table C.9 reports the result: the preferred coefficient moves from 0.802 to 0.798 (SE 0.052) and the R^2 from 0.469 to 0.468. The own-household overlap is immaterial to the validation.

Table C.9: Return Validation with Leave-One-Out Crop Means

	(1) No FE	(2) Year FE	(3) Crop-Set FE	(4) Crop-Set + Year FE
Expected return/HA, leave-one-out	1.559***	0.798***	1.534***	0.744***
	(0.058)	(0.052)	(0.059)	(0.052)
Full-sample coefficient (Table B.10)	1.563	0.802	1.539	0.748
N	7,084	7,084	7,084	7,084
R^2	0.109	0.468	0.114	0.476
p (H_0: $\beta = 1$)	0.000	0.000	0.000	0.000
Crop-Set FE	No	No	Yes	Yes
Year FE	No	Yes	No	Yes

Notes: The four specifications of Table B.10 on the same 7,084 pure-control household-years (rounds 0 to 2), with each household's expected return rebuilt from crop means that exclude its own farmer-crop mean. The "full-sample coefficient" row repeats Table B.10's estimates on this sample for comparison. Standard errors clustered at the household level in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

C.10: Yield, Price and Covariance Shares of Crop Risk

Section 2.2 describes two sources of systematic risk, weather acting on yields and prices acting on the value of what is harvested, and a natural-insurance channel through which the two offset. This section measures their shares in the crop-level risk that the covariance matrix prices, because the equilibrium reading of Figure A.8 in Section 5.1 rests on variance being mostly a property of the crop rather than of its price. Revenue per hectare is yield times price, so for a farmer-crop cell $\log(\text{revenue/HA}) = \log(\text{kg/HA}) + \log(\text{price per kg})$; demeaning within the cell as $\mathbf{C}$ is, $\mathrm{Var}(r) = \mathrm{Var}(y) + \mathrm{Var}(q) + 2\,\mathrm{Cov}(y, q)$, and the three terms over $\mathrm{Var}(r)$ are the yield, price and covariance shares. Revenue rather than profit is decomposed, because costs are a farmer-year object with no yield or price reading, and the exercise is support for an interpretation rather than a result the paper rests on. The closest published analogue is El Benni and Finger (2014), who decompose net revenue variance for Swiss crops at the farm level, in levels rather than logs and with cost terms carried

alongside price and yield. They find the price-yield covariance, the "natural hedge" of that literature, cuts net revenue variance by six percent for wheat and by as much as a third for sugar beet, an order of magnitude more than the offset found here; they also find costs contribute little, which is the same reason revenue rather than profit is decomposed above.

Quantities in kilograms come from the raw plot file and are matched to 97 percent of the cleaned pure-control plot-crop harvests; the price is the implied one, harvest value over quantity, so it is whatever rule valued the harvest. That rule is close to the community-round-crop median described in Section 2.2: the implied price is a single number in 77 percent of community-round-crop cells, and the shares below are computed on the prices actually applied. Because a log decomposition cannot carry a zero harvest, the table reports the share of each crop's plot-crop observations with zero harvest, which sit entirely in the yield term. The second block restricts to the 72 percent of harvests priced at their own community's median rather than the sample-wide fallback, which compresses price variance; the two blocks agree.

Table C.10 reports the result for the 14 crops with at least 100 farmer-crop cells (maize and groundnut at the species level, since yield and price are species objects and the high/low split is a cost split). Yield variation is 0.98 of the within-farmer variance of log revenue per hectare on an observation-weighted basis, running from 0.92 to 1.00 across crops; price variation is 0.04. The covariance share is -0.02 and negative for 11 of the 14 crops, the sign the natural-insurance argument of Newbery and Stiglitz (1981) predicts, but small: local prices offset only a few percent of yield variation, which is what Allen and Atkin (2022) would lead one to expect of rural markets that are partially integrated. Cotton is the exception with a positive covariance, consistent with a price set through the parastatal marketing chain rather than by local supply. Conditional on price, then, the risk a crop carries is a property of the crop and the weather, which is the premise the equilibrium reading of the risk-return line needs.

Table C.10: Yield, Price and Covariance Shares of Within-Farmer Revenue Variance, by Crop

		Zero	All harvests				Own-community priced			
Crop	Cells	harv. (%)	Yield	Price	2 Cov	ρ	Cells	Yield	Price	2 Cov
Groundnut	2,220	3	0.99	0.05	-0.04	-0.08	2,185	0.99	0.05	-0.04
Maize	2,210	1	0.98	0.05	-0.03	-0.07	1,481	0.98	0.05	-0.03
Rice	1,807	8	1.00	0.01	-0.02	-0.07	1,582	1.01	0.01	-0.02
Sorghum	1,663	4	1.00	0.02	-0.02	-0.07	594	1.00	0.02	-0.01
Cotton	1,432	3	0.92	0.06	0.02	0.05	1,349	0.92	0.06	0.03
Beans	769	18	0.98	0.03	-0.01	-0.03	428	0.98	0.03	-0.01
Pearl millet	760	3	1.00	0.05	-0.05	-0.10	270	0.96	0.06	-0.02
Okra	685	10	0.97	0.03	0.00	-0.01	373	0.98	0.03	-0.01
Other	505	47	0.97	0.01	0.02	0.13	55	0.99	0.01	0.00
Fonio	360	4	0.94	0.07	-0.01	-0.01	270	0.97	0.04	-0.01
Sweet potato	343	9	1.00	0.00	0.00	0.00	90	1.00	0.00	0.00
Sesame	147	11	1.00	0.00	0.00	-0.04	140	1.00	0.00	0.00
Yam	145	5	0.99	0.01	0.00	0.00	106	0.99	0.02	-0.01
Bambara groundnut	103	15	0.92	0.05	0.03	0.06	7	1.03	0.05	-0.08
Weighted mean			0.98	0.04	-0.02			0.98	0.04	-0.02

Notes: Pure-control plot-crop harvests with positive quantity and value (62,821 observations). For each farmer-crop cell with at least two seasons, log revenue per hectare, log yield (kg/HA) and log implied price (value/kg) are demeaned within the cell; the columns are $\mathrm{Var}(y)/\mathrm{Var}(r)$, $\mathrm{Var}(q)/\mathrm{Var}(r)$ and $2\,\mathrm{Cov}(y,q)/\mathrm{Var}(r)$, which sum to one, and ρ is the correlation of the yield and price deviations. "Cells" counts farmer-crop cells; "Zero harv." is the share of the crop's plot-crop observations with zero harvest, excluded from the logs. The right-hand block restricts to harvests in community-round-crop cells where at least one household sold, so the price is the community's own median rather than the sample-wide fallback. Maize and groundnut are at the species level. Crops with fewer than 100 cells are omitted from the table but carried in the macro file. The weighted mean weights crops by observations.

C.11: What the Return Validation Does and Does Not Test

Table B.10 reports the return validation; this section reads its R^2 pattern, which was moved here from the main text because it explains the construct rather than setting up the experiment. The R^2 pattern is informative about what the expected-return construct is and is not designed to do. By construction, $\bar{R}_c$ is a single number for each crop c, averaged across farmers and across years: it recovers the typical profit per hectare associated with choosing crop c rather than another, and it carries no community, household, or year dimension of its own. When $\bar{R}_c$ is regressed on realized profit without fixed effects, the residual contains everything that the averaging removed: the season, the local weather and price realizations behind it, and household-level idiosyncratic variation in cultivation practices. A low R^2 in the no-fixed-effects column is therefore the expected pattern rather than an indictment of the moment estimate. The construct was not designed to forecast realized profit at the household-year level, and a high R^2 in this column would imply that the cross-crop expected-return structure carried community-year information that, by definition, it does not.

Year fixed effects absorb the aggregate part of that residual, and the R^2 moves from 0.109 to 0.469. Crop-set fixed effects add little on their own, 0.115 against 0.109 without them: the 18 crop sets record which crops a household could have planted, which is a weak predictor of how any particular season turned out.

I use crop-set rather than community fixed effects here so that the validation exercise absorbs variation at the same level as the outcome variables I employ later, which are measured against crop-set specific efficient frontiers. The coefficient on $\bar{R}_c$ stays positive and precisely estimated in all four columns. It falls from 1.56 to 0.80 once year effects enter, because in the raw cross-section the construct partly stands in for whether a season was good; within a season, a one-unit increase in constructed expected return maps to 0.80 units of realized profit per hectare. Under the maintained assumption that crops carry common moments, expected portfolio return is $\mathbf{w}_h'\bar{\vec{R}}$, so the within-season slope implied by the framework is one rather than merely positive. The preferred column rejects that null ($p < 0.001$), and a slope of 0.80 is what classical attenuation in $\bar{R}_c$ would produce: the construct spreads returns across crops by roughly a quarter more than realized profits do. Read that way, the return dimension of the frontier is stretched relative to the data, so the dollar magnitudes I read off it in Section 7 should be treated as upper bounds on the return side. The sign and precision remain the relevant evidence that the cross-crop return structure recovered from the data is real, and the same cross-crop structure underlies the outcome variables in the treatment-effect analysis. The saturated column, which adds crop-set effects, is a robustness check on that reading; the preferred column and the treatment-effect specifications in Section 6 both use year fixed effects only. The risk construct cannot be validated this sharply, because a single season says nothing about dispersion and the comparison must therefore be made across groups of households rather than household by household; at that weaker standard, households the framework assigns more risk do bear more realized variation in profit per hectare (Appendix C.7). The expected-return moments are therefore valuable as descriptors of the inherent risk-return profile of each portfolio, even though their cross-sectional predictive power for individual household-year realized profit is, appropriately, modest.

Appendix D: Credibility of the Covariance Matrix

A natural concern with estimating crop return moments from a four-year panel is whether the resulting covariance matrix is stable across different realizations of weather and price states. If estimated correlations are highly sensitive to which years are observed, then portfolio properties constructed from these moments may be unreliable. This appendix collects the evidence that the covariance matrix $\mathbf{C}$ is a credible object, in four parts. Section D.1 uses the seven-wave Uganda LSMS-ISA panel to ask how far a four-wave estimate of the crop correlations sits from a longer-run one, pooled and within each of Uganda's regions. Section D.2 turns to the frontier built on that matrix: the Markowitz problem by which stable inputs can still produce unstable frontier corners, why the within-experiment comparison does not depend on a window-stable frontier, and a Ledoit-Wolf shrinkage exercise showing that the portfolio-SD treatment effect is unmoved by regularizing $\mathbf{C}$. Section D.3 re-estimates $\mathbf{C}$ from community-year rather than farmer-year deviations, as the robustness check on the revenue-share area allocation. Section D.4 re-estimates the correlations from cross-plot pairs alone, removing the observations on which the allocation convention makes two crops mechanically identical.

D.1: Ugandan LSMS-ISA Exercise

I estimate the covariance matrix from all seven waves of available data and then re-estimate it from each of the $\binom{7}{4} = 35$ possible four-wave subsets, treating the seven-wave estimate as the benchmark a four-year window is trying to recover. Each window mirrors the Mali sample on both margins that matter: a farmer enters a window only if observed in every one of its four waves, the analogue of Mali's balanced panel, and a crop enters a farmer's deviations only if that farmer grew it in at least two of those waves, as in Section 2.2. The benchmark pools every farmer observed in four or more waves, which is the union of the window samples, since a farmer with four waves is balanced in at least one window.

Pooling the entirety of the Ugandan LSMS-ISA into one national sample to make one national variance-covariance matrix fails to account for the fact that the data spans 4 administrative regions with diferring crop menus and agro-ecological zones present. This is not a problem I had to overcome with the Malian data as Sikasso is a single region with one menu. Because of this I am able to expand the exercise so that every four-wave window then yields 4 region-specific matrices, 140 in all, and four replications on samples of one region each (775 to 941 farmers, against 3,537 pooled).

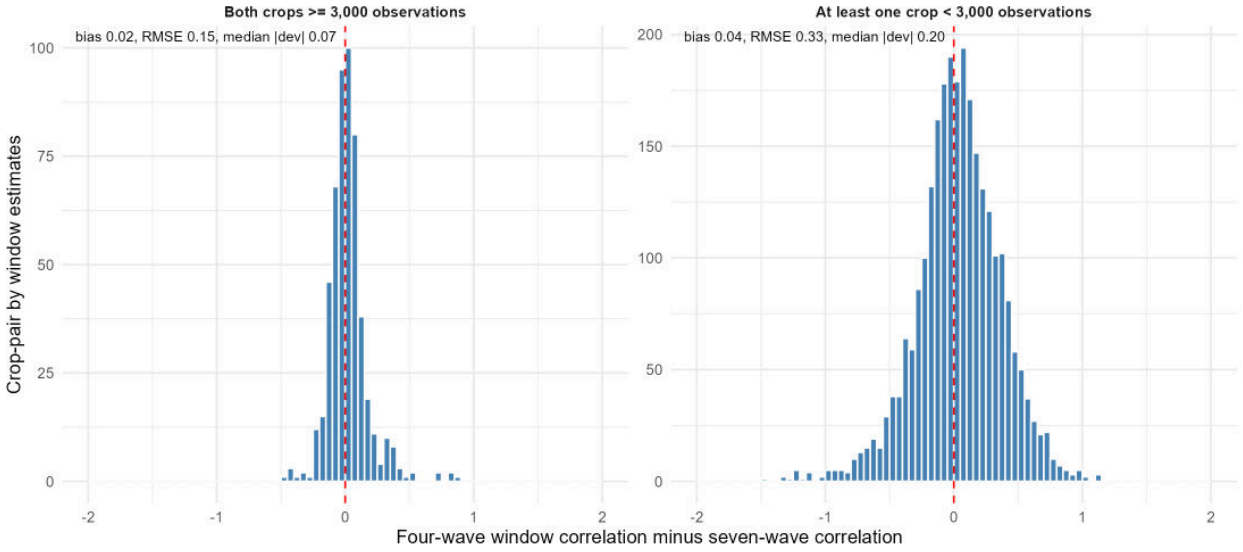
Figure D.1 plots, for every crop pair and every window, the window correlation minus the seven-wave correlation, split by whether both crops clear 3,000 observations. Both distributions are centered on the benchmark: the mean deviation is 0.02 among the 15 widely cultivated pairs and 0.04 over all 91 pairs. What differs is the spread. For widely cultivated pairs the root mean squared deviation is 0.15, 65% of

window estimates lie within 0.1 of the seven-wave value and 88% within 0.2; where either crop is rare the root mean squared deviation is 0.33 and only 51% lie within 0.2. The widely cultivated crops are the ones that dominate household portfolios and drive the treatment effects in Mali; the niche combinations that scatter are the ones that carry little weight.

Table D.1 reports the same deviation statistics region by region, with the widely cultivated set fixed at the national threshold so the split matches Figure D.1. The pattern replicates in every region: four-wave estimates of the widely cultivated pairs are centered on the region's seven-wave value (no mean deviation larger than 0.03 in absolute value), their root mean squared deviation runs from 0.18 to 0.25 against 0.15 pooled, and 65% to 79% of window estimates fall within 0.2 of it. In each of the 3 regions whose common crop set extends beyond the widely cultivated ones, the remaining pairs scatter more (root mean squared deviation 0.28 to 0.33). The modest rise in dispersion relative to the pooled sample is what a quarter of the farmers implies; the ordering is what transfers to Mali. The stable part of the covariance matrix is the same part in every region, so it is a property of widely cultivated crops in a four-wave window rather than of any one sample.

This exercise suggests that the four-year panel in Mali provides reasonably stable moment estimates for the crops that matter most for portfolio construction; but correlations involving rarely cultivated crops should be read with caution, and Section D.2 below takes up how that caution reaches the frontier.

Figure D.1: Four-Wave Window Correlations Relative to the Seven-Wave Estimate (Uganda LSMS-ISA)



Notes: Each observation is one crop pair in one four-wave window: the correlation estimated from that window minus the correlation estimated from all seven waves, for the 91 pairs among the crops estimable in every window (3,185 pair-window estimates). Left: both crops have at least 3,000 household-crop observations (15 pairs). Right: at least one crop has fewer (76 pairs). A window uses only farmers present in all four of its waves, and a crop enters a farmer's deviations only if grown in at least two of them; the seven-wave benchmark uses every farmer observed in four or more waves. Each window is four-sevenths of the benchmark, so deviations are shrunk toward zero by construction, and the 35 windows overlap, so the histograms describe dispersion rather than a sampling distribution. Dashed line at zero.

Table D.1: Within-Region Replication of the Four-Wave Window Exercise (Uganda LSMS-ISA)

			Widely cultivated pairs				Remaining pairs	
Region	Farmers	Crops	Pairs	Bias	RMSE	Within 0.2 (%)	Pairs	RMSE
Central	775	6	15	-0.03	0.18	79	—	—
Eastern	929	8	15	0.03	0.21	71	13	0.33
Northern	905	8	10	0.03	0.25	65	18	0.28
Western	941	8	15	0.02	0.21	68	13	0.30
Pooled, all regions	3,537	14	15	0.02	0.15	88	76	0.33

Notes: Each row repeats the exercise of Figure D.1 within one administrative region: 35 four-wave windows, each using only farmers present in all four of its waves, compared with the region's seven-wave estimate on every farmer observed in four or more waves. Crops is the number of crops estimable in every window of that region; pairs are counted among them. Widely cultivated pairs are those where both crops have at least 3,000 household-crop observations in the national sample, the same split as Figure D.1; remaining pairs are all others. Bias is the mean of window minus seven-wave correlation, RMSE its root mean square, and Within 0.2 the share of window estimates within 0.2 of the seven-wave value. The pooled row applies the same method to all regions together. Regions with fewer than 200 farmers are excluded.

D.2: Markowitz Error-Maximization and Sensitivity to Covariance Shrinkage

The frontier portfolios themselves, the global minimum-variance and maximum-Sharpe corners of the upper branch, can swing substantially across the same four-wave subsets even when the inputs are stable: built on the 6 widely cultivated crops alone, the minimum-variance σ has a coefficient of variation of 0.186 across the 35

windows and the maximum Sharpe ratio 0.803. This is the well-documented Markowitz error-maximization problem (Michaud, 1989; Best and Grauer, 1991; Ledoit and Wolf, 2004): the inverse $\mathbf{C}^{-1}$ that anchors $w_{\text{GMV}} \propto \mathbf{C}^{-1}\mathbf{1}$ and $w_{\text{max-Sharpe}} \propto \mathbf{C}^{-1}\boldsymbol{\mu}$ amplifies noise in the directions of $\mathbf{C}$'s smallest eigenvalues, so finite-sample noise gets channeled disproportionately into the frontier portfolio weights regardless of how stable the individual correlations and SDs are. The standard remedy in the finance literature when the number of assets is large relative to the time-series sample is Ledoit and Wolf (2004) shrinkage of $\mathbf{C}$ toward a low-parameter target.

The within-experiment treatment effect identification does not require a sample-window-stable frontier; it requires a *consistently-applied* frontier across treatment and control households, which the same-window comparison guarantees by construction. The relevant robustness exercise is therefore whether the headline portfolio-SD effect survives regularizing $\mathbf{C}$ in the way the Markowitz problem motivates. This subsection re-estimates $\mathbf{C}$ as a convex combination

$$\mathbf{C}_\alpha = \alpha\,\mathbf{C}_{\text{target}} + (1-\alpha)\,\mathbf{C}_{\text{sample}}, \qquad \alpha \in \{0, 0.25, 0.50, 0.75, 1.00\},$$

where $\mathbf{C}_{\text{target}}$ is the constant-correlation target with off-diagonal entries set to the sample mean correlation $\bar{\rho} = 0.292$, recomputes household-level portfolio standard deviation under each shrunk covariance matrix using the existing acre-share weights, and re-estimates the grant's treatment effect on PortSD. Results are reported in Table D.2.

Table D.2: Treatment Effect on Portfolio SD under Ledoit-Wolf Shrinkage of $\mathbf{C}$

Shrinkage α	Coef.	SE	t-stat	p-value
0.00	-4.011**	(1.832)	-2.19	0.0286
0.25	-4.104**	(1.839)	-2.23	0.0257
0.50	-4.197**	(1.852)	-2.27	0.0235
0.75	-4.289**	(1.870)	-2.29	0.0219
1.00	-4.382**	(1.894)	-2.31	0.0208
N = 6,276; sample $\bar{\rho} = 0.292$; range $[\rho_{\min}, \rho_{\max}] = [-0.179, 0.560]$				

Notes: The covariance matrix $\mathbf{C}_\alpha = \alpha\mathbf{C}_{\text{target}} + (1-\alpha)\mathbf{C}_{\text{sample}}$ shrinks the farmer-level sample covariance toward a constant-correlation target with $\bar{\rho} = 0.292$ at intensity α. Each row reports the grant treatment effect on Portfolio SD when household portfolio SDs are recomputed under $\mathbf{C}_\alpha$, using the headline ANCOVA specification of Section 6: the Beaman et al. (2023) controls, round indicators, a baseline control for the dependent variable, and standard errors clustered at the household. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The treatment effect is essentially unchanged across the full shrinkage range. The point estimate moves from -4.011 at $\alpha = 0$ (no shrinkage, sample covariance) to -4.382 at $\alpha = 1$ (full shrinkage to the constant-correlation target), a movement of well under one standard error, and statistical significance is preserved at the 5% level throughout ($p = 0.029$ at $\alpha = 0$, $p = 0.021$ at $\alpha = 1$). The headline portfolio-SD treatment effect

is therefore not contingent on the particular sample covariance estimate: the variance reduction documented in the main text reflects a treatment-induced shift in portfolio composition rather than fragile dependence on the off-diagonal entries of $\mathbf{C}$ that are most exposed to thin-pairwise-sample sampling error, even though the frontier itself is sensitive to the regularization choice.

D.3: Community-Level Aggregation as an Area-Allocation Robustness

The headline covariance matrix $\mathbf{C}$ is estimated at the farmer-year level, with intercropped-plot areas allocated to constituent crops by revenue share (Section 2.2). A natural concern with revenue-share allocation is yield-realization endogeneity: when a crop fails on an intercropped plot, its revenue share collapses, and so does the allocated area in the denominator of π_c/A_c. The allocator co-varies with the realized state of the world.

Three alternative allocation rules were considered as direct remedies but each carries its own bias. Equal-split (1/N) imposes an empirically unreasonable assumption that intercropped plots have equal physical allocation across crops, and mechanically amplifies measurement noise via an $N \times \text{rev}_c$ term that destroys statistical power. Seed-cost-share is infeasible because only 19–32% of plot-crop observations in a given round record any seed expenditure (the rest used home-saved seed, with no recorded value). Sole-cropped-only $\mathbf{C}$ avoids allocation entirely by restricting to the sole-cropped subset, but the off-diagonal pairwise samples become very thin (a farmer must sole-crop both crops in the same year), and the resulting matrix requires PSD-fixing because the off-diagonal estimates are ill-conditioned.

The robustness reported here aggregates to (community, year) before estimating $\mathbf{C}$. For each (community, year, crop), I average π/HA across pure-control farmers cultivating that crop (sole-cropped or intercropped, with intercropped allocation handled by the same revenue-share rule as the headline). I then demean by (community, crop) across years and compute $\mathbf{C}_{\text{comm}}$ from the pooled (community, year) deviations. This procedure has three properties that make it the cleanest robustness available for the area-allocation question. First, idiosyncratic farmer-specific noise, including any farmer-level allocation idiosyncrasies, washes out in the community-year average. Second, pairwise overlap is dense: most communities have at least some farmer cultivating each major crop in any given year, so the (community, year) panel covers nearly every crop pair without selection. Third, the resulting matrix is positive-semi-definite without forcing, indicating well-conditioned covariance estimates.

Table D.3 reports the headline portfolio-SD treatment effect under both the farmer-level $\mathbf{C}$ (the headline) and the community-level $\mathbf{C}$. The point estimates remain close: −\$4.01 per hectare under farmer-level estimation versus −\$3.62 per hectare under community-level estimation, a gap of 0.21 farmer-level standard errors. The community-level standard error is in fact tighter (1.55 versus 1.83), reflecting the noise reduction

from cross-farmer averaging. The farmer-level estimate has $p = 0.029$ and the community-level estimate $p = 0.020$.

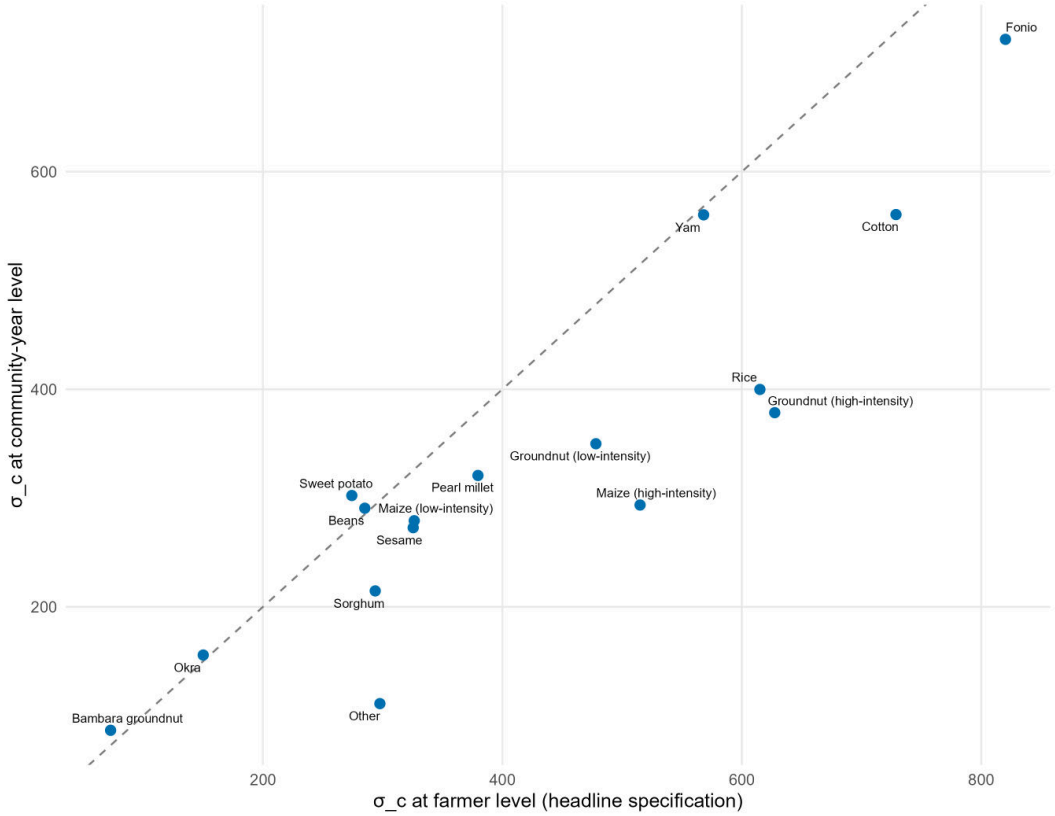
Table D.3: Treatment effect on portfolio SD: farmer-level vs community-level **C**

	Farmer-level **C** (headline)	Community-level **C**
Grant TE on Portfolio SD	−4.01**	−3.62**
	(1.83)	(1.55)
p-value	0.0286	0.0196
Control mean ($/HA)	316	269
TE as % of control mean	−1.27%	−1.34%

Notes: Both columns use the headline ANCOVA specification of Section 6: the Beaman et al. (2023) controls, round indicators, a baseline control for the dependent variable, and standard errors clustered at the household. The community-level column substitutes a covariance matrix estimated from (community, year) deviations of crop-average π/HA for the headline farmer-level **C**, holding portfolio weights and treatment-effect specification fixed. * p<0.10, ** p<0.05, *** p<0.01.

The crop-level standard deviations themselves also align closely: the Spearman rank correlation between farmer-level and community-level σ_c rankings is 0.87, and the Pearson correlation of the σ_c levels is 0.90 (Figure D.2). The mean ratio $\sigma_{\text{comm}}/\sigma_{\text{farmer}}$ is 0.82, reflecting the noise reduction from cross-farmer averaging. Community-level σ_c captures only the systematic component of crop-return variance, missing the idiosyncratic component that farmer-level **C** also reflects. Major crops are stable across estimation strategies (cotton 15→15, fonio 16→16, rice 13→13, sorghum 5→4); modest reorderings appear for crops where idiosyncratic farmer-level variance is a larger share of total (σ_{other} drops toward the safe end of the distribution).

Figure D.2: Crop-level SD: farmer-level vs community-level estimation



Notes: Each point is a crop. The horizontal axis is σ_c from the farmer-level $\mathbf{C}$ used as the headline; the vertical axis is σ_c from the community-level $\mathbf{C}$ in this appendix. The diagonal reference line marks perfect agreement. Spearman rank correlation 0.87; Pearson correlation of levels 0.90.

The community-level $\mathbf{C}$ measures a different object than the farmer-level $\mathbf{C}$: it captures community-systematic crop covariation only, ignoring the idiosyncratic component that individual farmers also experience. This is consistent with the systematic-vs-idiosyncratic decomposition reported in Appendix G.3, where roughly 85% of farmer-level portfolio variance is community-systematic. For the area-allocation robustness question, the relevant comparison is between the headline TE under farmer-level Σ and the TE under any cleaner-estimated alternative. The community-level estimate provides such an alternative, with dense pairwise coverage, well-conditioned moment estimates, and a TE that matches the headline within rounding. The headline portfolio-SD reduction is therefore robust to a fundamentally different estimation strategy that uses different aggregation, different effective sample, and different bias structure.

D.4: Correlations Estimated from Cross-Plot Pairs Only

Appendix C.1 notes that allocating both area and input costs by revenue share gives every crop on an intercropped plot the same profit per hectare, so the within-plot residuals of two co-located crops are mechanically

perfectly correlated. That appendix answers the concern with prevalence, reporting how few farmer-year crop pairs are co-located. This section replaces the prevalence argument with a direct test: re-estimate every correlation using only the farmer-years in which that pair did *not* share a plot, then rebuild the portfolio risk outcome and re-run the headline specification.

The restriction is targeted rather than blunt. Restricting to sole-cropped plots, discussed in Appendix D.3, discards an observation whenever *either* crop shares a plot with anything at all, which is why its off-diagonal samples go thin and the resulting matrix has to be forced positive semi-definite. Here the crop variances are left exactly as the headline estimates them, since a variance involves no pair, and only the correlations are re-estimated. For each pair, the farmer-years dropped are those in which that specific pair sat on a shared plot; a crop intercropped with something else remains usable for every pair it is not co-located in. That retains 77.0% of pairwise observations rather than the small fraction the sole-cropped restriction leaves, with a median of 378 farmer-years behind each correlation.

The correlation structure barely moves. Across the off-diagonal entries the mean absolute change is 0.028 and the largest is 0.170, and the rebuilt matrix is positive semi-definite without any eigenvalue forcing, which the sole-cropped alternative is not. Table D.4 reports the consequence for the result that runs through $\mathbf{C}$. The grant's effect on portfolio standard deviation is -4.13 USD per hectare under cross-plot correlations against -4.01 in the headline, with $p = 0.026$ against 0.029. The effect is marginally larger in magnitude once the mechanically correlated observations are removed, not smaller. Whatever the revenue-share convention writes into the off-diagonals, it is not what produces the measured risk reduction; if anything it works against it.

Table D.4: Portfolio-SD Treatment Effect under Cross-Plot-Only Correlations

	Effect	SE	p	Control mean	% of control
All farmer-year pairs (headline)	-4.01	(1.83)	0.029	316	-1.27
Cross-plot pairs only	-4.13	(1.86)	0.026	313	-1.32

Notes: Both rows use the headline ANCOVA specification of Section 6 and the same household portfolio weights; they differ only in the matrix $\mathbf{C}$ that prices those weights. The first row is the headline matrix. The second re-estimates each off-diagonal correlation after dropping the farmer-years in which that crop pair shared a plot, leaving the crop variances unchanged, and retains 77.0% of pairwise observations. Standard errors clustered at the household. Control means differ across rows because the rebuilt matrix changes the level of the outcome for every household, not only for treated ones.

Appendix E: Randomization Inference and Multiple Hypothesis Testing

To account for the large number of outcome variables tested simultaneously and to provide distribution-free inference, I supplement the analytical standard errors reported in the main tables with randomization inference (RI) using 10,000 permuted treatment assignments from the original experimental design of Beaman et al. (2023).

For each permutation, I replace the actual grant indicator with the permuted assignment and re-estimate the full ANCOVA specification on all 16 outcomes. Following the randomization-inference implementation in the Beaman et al. (2023) replication package, the sample restrictions that define the estimation sample are re-applied within each permuted assignment, so households outside the relevant treatment cells under a given draw are excluded for that draw; the estimation sample therefore varies across permutations exactly as the realized sample varies with the realized assignment, keeping the reference distribution faithful to the assignment mechanism. The individual RI p-value for each outcome is the fraction of permutations yielding a t-statistic at least as large in absolute value as the actual t-statistic. To correct for multiple testing within outcome families, I apply the Westfall-Young stepdown procedure (Westfall and Young, 1993), which controls the familywise error rate while accounting for the joint dependence structure of the test statistics. I define four outcome families: (i) ex-ante cultivation decisions, (ii) ex-post realized outcomes, (iii) MPT-constructed portfolio measures, and (iv) efficiency measures. The Westfall-Young adjusted p-values are reported as a row in Tables B.13 and B.14; the efficiency family also contains origin distance, which is adjusted with the rest and reported in Appendix H.

A potential concern with the MPT-based outcomes is that they are constructed from estimated crop return moments (expected returns, variances, and covariances), introducing a generated regressor problem: the analytical standard errors treat these constructed outcomes as observed data and may understate uncertainty. A standard remedy would be to bootstrap the full estimation chain (resampling farmers, re-estimating crop moments, re-constructing portfolio outcomes, and re-running the treatment regressions) to propagate first-stage estimation error into the standard errors. However, randomization inference largely obviates this need. The RI p-values do not depend on the standard errors at all; they compare the actual test statistic against the permutation distribution under the sharp null of no treatment effect for any unit. Whether the analytical standard errors are too small or too large is irrelevant to the RI p-value, which is exact under the null by construction. Additionally, the crop return moments are estimated exclusively from pure control households (a separate sample from the treatment comparison), which limits the scope for first-stage estimation error to contaminate the treatment effect estimates. For these reasons, I rely on randomization

inference as the primary basis for statistical conclusions about the MPT-constructed outcomes.

Appendix F: Heterogeneous Treatment Effects

Table F.1 examines whether the grant's portfolio effects vary by baseline household characteristics. Each panel interacts the grant indicator with a standardized (mean zero, unit standard deviation) baseline moderator: portfolio standard deviation, distance to the efficient frontier, agricultural asset holdings, and number of crops cultivated. The interaction coefficient represents the additional treatment effect per one standard deviation increase in the moderator. Because these interactions are linear, they test only for a monotone gradient and cannot capture the non-monotone, both-tails heterogeneity documented on the wealth index in Section 6.3; the two analyses are complementary, not conflicting. None of the 20 interactions are significant at the 10% level. That null follows from the tercile pattern of Section 6.3 rather than contradicting it. A linear interaction tests for a *monotone* moderator slope, whereas the tercile analysis shows the grant's effect on portfolio risk is *non-monotone*: both the poorest and richest terciles reduce risk while the middle does not. A pattern that moves in the same direction at both tails and flattens in between has a linear slope near zero, so a single grant-by-moderator interaction is designed to miss it: the opposite-tail movements average out. The two results are therefore complementary rather than contradictory: the linear interactions rule out a simple monotone gradient in these moderators, while the tercile splits recover the non-monotone, both-tails structure that no single interaction coefficient can represent. As a further check, I use the authors' household assets index as the heterogeneity dimension in place of the agricultural asset index reported in Table F.1. Baseline portfolio risk is essentially uncorrelated with household assets ($r = -0.005$), and the treatment effect on portfolio SD does not vary with baseline wealth (interaction coefficient $= 0.98$, $p = 0.583$).

Table F.1: Heterogeneous Treatment Effects: Grant × Baseline Characteristics

	Portfolio SD	E[Profit]/HA	Sharpe Ratio	EF Distance	Div. Ratio
	(1)	(2)	(3)	(4)	(5)
Moderator: Baseline Portfolio SD					
Grant	-3.89**	-1.57	0.70***	-1.73***	1.19***
	(1.79)	(1.71)	(0.26)	(0.52)	(0.34)
Grant × Z	-1.85	-1.20	-0.15	0.26	-0.21
	(2.41)	(1.91)	(0.33)	(0.67)	(0.40)
Moderator: Baseline EF Distance					
Grant	-3.86**	-1.81	0.67***	-1.48***	1.16***
	(1.81)	(1.73)	(0.26)	(0.53)	(0.34)
Grant × Z	1.63	1.72	0.03	-0.03	-0.12
	(2.45)	(2.09)	(0.38)	(0.77)	(0.38)
Moderator: Baseline Ag Assets					
Grant	-4.01**	-2.37	0.67***	-1.49***	1.15***
	(1.83)	(1.76)	(0.26)	(0.53)	(0.35)
Grant × Z	0.52	-1.00	-0.45	0.63	0.30
	(1.95)	(1.81)	(0.27)	(0.58)	(0.35)
Moderator: Baseline N Crops					
Grant	-4.30**	-2.49	0.64**	-1.41***	1.13***
	(1.82)	(1.71)	(0.26)	(0.53)	(0.35)
Grant × Z	0.18	-0.08	-0.11	0.46	-0.06
	(2.07)	(1.87)	(0.28)	(0.56)	(0.36)
N	6,276	6,276	6,276	6,276	6,276

Notes: Each panel interacts the grant indicator with a standardized (mean 0, SD 1) baseline characteristic. Columns report results for six MPT and efficiency outcomes. All specifications include the baseline value of the dependent variable, the authors' preferred control suite, and a year fixed effect. Standard errors clustered at household level in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Appendix G: Robustness to Alternative Specifications

G.1: Cotton-Excluded Robustness

The cultivator-pool spread analysis in Section 2.4 (and the supporting figure in Appendix J) flags cotton as the major crop whose cultivators are visibly skewed by farm size: only 5.8% of cotton cultivators come from the smallest baseline-farmland quintile, compared to 11–19% for every other major crop. This concentration is consistent with the labor and processing requirements of cotton production rather than with capability-based risk sorting, but it raises a natural robustness question: are the headline portfolio treatment effects driven by the cotton subsample where Q1 representation is poor?

Table G.1 reruns the full set of headline outcomes restricting the analysis sample to households who did not cultivate any cotton at baseline (year 2010). The cotton-excluded sample is roughly half the size of the full sample, which inflates standard errors mechanically. With that caveat, the grant effect on portfolio standard deviation is −2.65 USD per hectare in the cotton-excluded sample against −4.01 in the full sample (Table G.1): the same sign, about two thirds of the magnitude, and no longer distinguishable from zero on half the sample. The treatment effect on the number of crops cultivated is also nearly identical (0.308 vs. 0.333), and the normalized HHI effect is in fact stronger in the cotton-excluded sample (−1.39 vs. −0.87).

Table G.1: Cotton-Excluded Robustness for Headline Treatment Effects

	Portfolio SD ($/HA)		Expected return ($/HA)		Sharpe ratio (×100)		N crops		Normalized HHI	
	Full	No-cotton	Full	No-cotton	Full	No-cotton	Full	No-cotton	Full	No-cotton
Grant	-4.011**	-2.648	-2.376	0.149	0.668**	0.956**	0.333***	0.308***	-0.872**	-1.388**
	(1.832)	(2.462)	(1.759)	(2.394)	(0.259)	(0.411)	(0.069)	(0.096)	(0.434)	(0.707)
N	6,276	3,382	6,276	3,382	6,276	3,382	6,276	3,382	6,276	3,382

Notes: The headline ANCOVA specification of Section 6: the Beaman et al. (2023) controls, round indicators, a baseline control for the dependent variable, and standard errors clustered at the household, comparing the full sample (treatment arms 4 and 5, rounds 1–2) to the subset of households who did not cultivate cotton at baseline (year 2010). SEs clustered at household level. * p<0.10, ** p<0.05, *** p<0.01.

The cultivator-selection concern raised in Section 2.4 therefore does affect but not entirely reverse the headline result: removing every cotton-cultivating household leaves the crop-count and concentration effects intact, the Sharpe ratio effect intact and the portfolio-SD point estimate of the same sign, though the smaller sample cannot establish it on its own.

G.2: Log-Linear Decomposition of Portfolio SD into Volatility, Concentration, and Correlation Channels

To complement the systematic-versus-idiosyncratic decomposition in Appendix G.3, this subsection reports a CCS-style multiplicative decomposition of portfolio standard deviation (Calvet et al., 2007):

$$\sigma_p \;=\; \sigma_a \cdot \sqrt{1/H} \cdot \delta,$$

$$\ln \sigma_p \;=\; \ln \sigma_a \;+\; \tfrac{1}{2} \ln(1/H) \;+\; \ln \delta,$$

where $\sigma_a = \sum_i w_i \sigma_i$ is the weight-averaged crop SD (the selection channel, capturing whether households select into cultivating inherently lower- or higher-volatility crops), $H = \sum_i w_i^2$ is the Herfindahl of weights so $1/H$ is the effective number of crops (the specialization channel), and δ is the residual diversification factor capturing correlation exploitation that is not summarized by the first two channels. The decomposition holds exactly at the farmer-year level.

Table G.2 reports treatment effects on the natural logarithm of each channel, estimated under the headline ANCOVA specification.

Table G.2: Treatment effects on log channels of the portfolio-SD decomposition

	ln Portfolio SD (total)	ln σ_a (volatility)	0.5· ln $1/H$ (concentration)	ln δ (correlation)
Grant	-0.0129**	-0.0051	0.0207***	-0.0296***
	(0.0055)	(0.0049)	(0.0061)	(0.0083)
% of total	100%	39%	-160%	229%
N	6,276	6,276	6,276	6,276

Notes: Each column is a separate regression under the headline ANCOVA specification of Section 6: the Beaman et al. (2023) controls, round indicators, a baseline control for the dependent variable, and standard errors clustered at the household. The selection channel (column 2) is the weighted-average individual crop SD; the specialization channel (column 3) is half the log of the effective number of crops $1/H$; the correlation channel (column 4) is the residual that closes the decomposition at the per-farmer-year level. Channel-level coefficients do not sum mechanically to the total because each regression uses its own baseline ANCOVA control. * p<0.10, ** p<0.05, *** p<0.01.

Table G.3: Treatment effects on the additive variance decomposition

	Portfolio variance (total)	Diagonal $\sum_i w_i^2\sigma_i^2$	Off-diagonal $2\sum_{i<j} w_i w_j \mathrm{Cov}_{ij}$
Grant	-2442*	-2678**	176
	(1302)	(1219)	(431)
Control mean	102,859	61,210	41,648
% of control mean	-2.4%	-4.4%	0.4%
N	6,276	6,276	6,276

Notes: Each column is a separate regression under the headline ANCOVA specification of Section 6: the Beaman et al. (2023) controls, round indicators, a baseline control for the dependent variable, and standard errors clustered at the household, matching Table G.2. The diagonal and off-diagonal columns sum exactly to the total: the identity is verified in code and holds to within 10^{-11}. The diagonal carries 60% of control-group portfolio variance. * p<0.10, ** p<0.05, *** p<0.01.

Before reading the channels, it helps to see the decomposition this one replaces. Table G.3 runs the headline ANCOVA on the additive variance identity $\sigma_p^2 = \sum_i w_i^2\sigma_i^2 + 2\sum_{i<j} w_i w_j \mathrm{Cov}_{ij}$, which holds exactly at the farmer-year level, so the two component columns sum to the total. Read naively, these results allude to the entirety of the variance reduction coming through a selection story on the diagonal of movement into safer crops. This reading ignores the squared terms on the weights that build in a specialization penalty in the formulation of the diagonal variance measure. I reference this again below in my discussion of Table G.2.

Table G.2's decomposition rests on the identity $\ln \sigma_p = \ln \sigma_a + \frac{1}{2}\ln N_{\text{eff}} + \ln \delta$, where $N_{\text{eff}} = 1/\sum_i w_i^2$ is the effective number of crops and δ is the residual diversification factor. All three terms enter positively, so a channel contributes to the headline reduction only when its own estimate is negative, and two of the three do. The correlation channel is the strongest: grant recipients hold crop combinations whose imperfect correlations pull portfolio risk below what breadth alone would deliver, with the contribution from $\ln \delta$ equal to -0.030 ($p < 0.001$). The volatility channel moves the same way but is weaker and imprecise: β on $\ln \sigma_a = -0.005$, $p = 0.30$, directionally consistent with a small shift toward lower-volatility crops on a weighted-average basis.

The breadth channel runs the other way, and mechanically so: grants raise the effective number of crops cultivated, and because N_{eff} enters positively, that spreading raises the breadth term by $+0.021$ ($p < 0.001$), consistent with the extensive-margin entry documented in Section 5.6. This coefficient is not a risk increase, in this parameterization the benefit of holding more crops is recorded in δ rather than in breadth, which is why the two channels carry opposite signs and why the correlation term more than offsets the breadth term.

The level-space decomposition in Table G.3 shows why it cannot answer that question. The whole reduction sits in the diagonal ($-2,678$, $p = 0.028$, 4.4% of its control mean) while the off-diagonal, the literal covariance term, is indistinguishable from zero ($+176$, $p = 0.68$). Read naively that says households selected safer crops and gained nothing from how those crops comove. But the diagonal $\sum_i w_i^2\sigma_i^2$ falls mechanically

whenever land spreads across more crops. This is because the squared weights shrink, and Section 5.6 shows the grant did exactly that. The diagonal therefore pools the volatility and concentration channels, and the level decomposition cannot separate a household that switched into safer crops from one that simply spread the same crops more evenly. This is the same squared-weight contamination that motivates the usage of the Selection Ratio in Section 2.6.

The two decompositions are not in conflict: what the level version books to the diagonal, the log version splits between breadth and δ. This is why δ should be read as a residual rather than as a pure correlation term, as it carries the diversification benefit of the added crops as well as any improvement in how they comove.

G.3: Systematic vs Idiosyncratic Variance Decomposition Against a Community Benchmark

To complement the log-linear channel decomposition in Appendix G.2, this subsection reports a decomposition of household portfolio variance into systematic (community-level) and idiosyncratic (household-specific) components in the spirit of Calvet, Campbell and Sodini (CCS) (Calvet et al., 2007, eq. 2). For each farmer h in community c, define

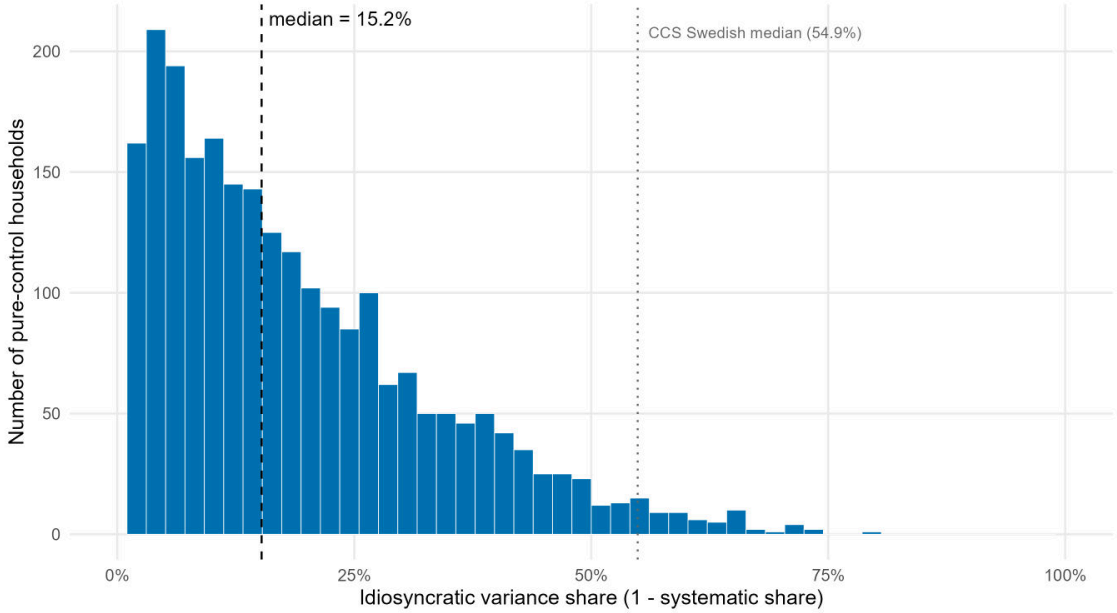
$$\sigma_h^2 \;=\; \beta_h^2\,\sigma_B^2 \;+\; \sigma_{i,h}^2,$$

where $\sigma_h^2 = w_h'\mathbf{C}w_h$ is the household's portfolio variance, $\sigma_B^2 = B_c'\mathbf{C}B_c$ is the variance of the community benchmark portfolio B_c, $\beta_h = w_h'\mathbf{C}B_c/\sigma_B^2$ is the regression beta of the household portfolio on the benchmark, and $\sigma_{i,h}^2$ is the idiosyncratic residual. The benchmark portfolio B_c is the hectare-weighted crop allocation of the entire community, computed from pure-control farmers at baseline (year 2010) so the benchmark is exogenous to treatment. This is the production-composition analogue of the value-weighted market portfolio that CCS use for Swedish stocks, and is the natural CAPM-style benchmark for a smallholder agricultural setting.

The headline statistic, the idiosyncratic variance share $\sigma_{i,h}^2/\sigma_h^2$, reveals a striking pattern (Figure G.1). Among baseline pure-control households, the median idiosyncratic variance share is 15.2%, with a mean of 18.9%; the median β_h is 0.99. For comparison, Calvet et al. (2007) report a median idiosyncratic share of 54.9% for Swedish equity investors against the value-weighted Swedish stock market. Mali smallholders' portfolio risk is therefore roughly 85% systematic, a much higher systematic share than CCS find for Swedish equity portfolios, and the opposite of the prior one might bring to a setting with thin crop menus and no inter-household pooling. The economic interpretation is that within each community, individual farmer portfolios cluster tightly around the community-aggregate composition because farmers grow largely the

same set of staples; the bulk of within-community portfolio variation in returns is therefore community-level (driven by weather, prices, and shared shocks) rather than idiosyncratic to the household.

Figure G.1: Distribution of idiosyncratic variance share at baseline



Notes: Idiosyncratic variance share is $\sigma^2_{i,h}/\sigma^2_h$ where σ^2_h is household h's portfolio variance and $\sigma^2_{i,h}$ is the residual after projecting on the community benchmark B_c. Sample: baseline pure-control households (year 2010). The dashed reference line marks the median (15.2%); the dotted reference line marks the CCS Swedish median (54.9%) reported in Calvet et al. (2007).

Table G.4 reports treatment effects on the four components: β_h, the idiosyncratic share, and the systematic and idiosyncratic variance levels. Grants leave β_h essentially where it was (a change of 0.005, $p = 0.366$) and reduce the idiosyncratic share by 0.013 ($p = 0.005$), 5.6% of its control mean. The variance-level decomposition puts the reduction on the idiosyncratic side: the idiosyncratic component falls by $1,544$ (6.5% of its control mean, $p = 0.024$), while the systematic component's estimate of $-1,100$ (1.4%, $p = 0.301$) is not distinguishable from zero.

Table G.4: Treatment effects on the systematic vs idiosyncratic decomposition

	β vs commune benchmark	Idiosyncratic share (%)	Idiosyncratic variance	Systematic variance
Grant	-0.005	-1.308***	-1544.8**	-1100.9
	(0.006)	(0.464)	(686.3)	(1064.1)
Control mean	0.867	23.4%	23760	79099
N	6,276	6,276	6,276	6,276

Notes: Each column is a separate regression under the headline ANCOVA specification of Section 6: the Beaman et al. (2023) controls, round indicators, a baseline control for the dependent variable, and standard errors clustered at the household. The benchmark portfolio B_c is the hectare-weighted community-level crop allocation computed from pure-control farmers at baseline. * p<0.10, ** p<0.05, *** p<0.01.

The substantive reading is that the grant moves households toward the community-average composition rather than away from community-level risk. Their exposure to community shocks, β_h, does not move, and the systematic component of their variance is not measurably lower; what falls is the idiosyncratic component, the part of portfolio variance that comes from deviating from the community's crop mix. This is what the cardinality-constrained mechanism of Section 5.6 predicts: relaxing the fixed-cost constraint lets farmers add the staple crops that make up the bulk of community production but that resource-constrained farmers could not afford to add at small shares, which pulls their portfolios toward the community-typical composition and reduces their idiosyncratic deviation from it.

The decomposition also qualifies how the risk reduction should be read against consumption. What the grant moves is the component that within-village transfers could in principle absorb, since community-wide shocks are common to a household's potential insurers and cannot be pooled locally, so the pass-through to consumption volatility suggested in Section 7 rests on the incompleteness of informal risk sharing in this agro-ecology (Kazianga and Udry, 2006) rather than on the size of the variance reduction alone.

This decomposition does not require a risk-free rate, which is why I do not report Sharpe-based versions of the same statistic, following the discussion in Section 2.1 and the rejection of cash-leasing as R_f in California crop markets in Blank (1991). The paper uses the variance-decomposition half of the CCS toolkit and avoids the Sharpe-based half.

G.4: Higher Moments of Crop Returns and the Mean-Variance Approximation

The mean-variance framework of Modern Portfolio Theory treats the first two moments of the return distribution as sufficient for portfolio-choice analysis, an approximation that is exact under quadratic utility or jointly Gaussian returns and is otherwise a Taylor expansion of expected CRRA utility around the portfolio

mean. This appendix documents the higher-moment properties of the crop-return distributions in this setting and asks whether the grant-induced portfolio shift moves households along a higher-moment dimension that the mean-variance bound fails to capture.

Crop-level higher moments. Skewness and excess kurtosis of farmer-demeaned crop returns are reported alongside the first two moments in Table B.6. All sixteen crops have positive skewness (range 0.4–4.6) and positive excess kurtosis (range 5.1–71.9), consistent with the production environment described in Section 3: returns are bounded below by drought-induced near-zero harvests and pulled to the right by occasional good-season or price-spike realizations, generating right-skewed distributions with heavier tails than the Gaussian benchmark. Cotton stands out with the heaviest tails (excess kurtosis 71.9), reflecting the cash crop's concentration of high-investment cultivation. The heterogeneity in higher moments across crops is therefore meaningful: crops in this menu differ not only in their mean and variance but also in the asymmetry and tail-thickness of their return distributions.

Treatment effect on portfolio-weighted higher moments. For each household-year, I construct portfolio-weighted analogues of the crop-level moments, $\text{PortSkew}_{i,t} = \sum_c w_{ic,t}\,\text{Skew}_c$ and $\text{PortExKurt}_{i,t} = \sum_c w_{ic,t}\,\text{ExKurt}_c$, where $w_{ic,t}$ is household i's acre share on crop c at year t and the crop-level moments are estimated as in Table B.6. These linear-in-weights aggregates do not give the true portfolio skewness or kurtosis (which would require coskewness and cokurtosis tensors that the four-year panel cannot identify); they answer a narrower descriptive question: are grant recipients tilting their land toward crops with more or less favorable higher-moment profiles?

Table G.5: Treatment Effect on Portfolio-Weighted Higher Moments

	(1) PortSkew	(2) PortExKurt
Grant (vs Control)	-0.011	-0.344
	(0.014)	(0.282)
Control mean	2.256	23.448
Treatment mean	2.250	23.344
N	6276	6276
R^2	0.142	0.324

Notes: ANCOVA specification matching the headline outcome regressions: each outcome is regressed on grant receipt, its own baseline value, an indicator for a missing baseline value, the authors' control suite and round indicators. Sample restricted to the authors' preferred sample, treatment arms 4 and 5, follow-up rounds 1–2; standard errors clustered at the household level. PortSkew and PortExKurt are land-share-weighted averages of crop-level skewness and excess kurtosis from Table B.6; they are not the true portfolio higher moments (which would require co-skewness and co-kurtosis tensors not estimable from a four-wave panel). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The treatment effects are reported in Table G.5, estimated by the same ANCOVA specification as the headline outcomes (Section 6). The point estimates on both PortSkew and PortExKurt are negative but in-

distinguishable from zero at conventional levels. Treatment recipients shift their land-share-weighted averages of skewness and excess kurtosis very slightly downward, but neither effect approaches statistical significance. The portfolio-aggregated higher-moment exposures of treated households are essentially unchanged by the grant is one conclusion, the fact that these land-share weighted averages are not holistically capturing the right concepts that would have been moved by treatment is an alternative explanation for what we see here.

Crop-level comparison: extensive-margin entries vs. commonly held crops. The portfolio-aggregated null obscures a substantively interesting pattern at the crop level. The three crops with significant positive treatment effects on the extensive margin in Figure A.25 (Bambara groundnut, high-intensity groundnut, pearl millet) have noticeably less extreme higher moments than the most commonly cultivated crops at baseline, as Table G.6 shows.

Table G.6: Higher Moments: Grant-Added Crops vs. Top-5 Commonly-Held Crops

Group	Crop	E[$/HA]	SD	Skew	Ex.Kurt
Commonly-held (top 5 baseline)	Cotton	531	729	3.70	71.88
Commonly-held (top 5 baseline)	Maize (high-intensity)	398	515	0.42	5.96
Commonly-held (top 5 baseline)	Groundnut (low-intensity)	373	478	2.27	13.12
Commonly-held (top 5 baseline)	Rice	361	615	2.27	11.58
Commonly-held (top 5 baseline)	Sorghum	232	294	1.73	9.17
Grant-added (extensive margin)	Groundnut (high-intensity)	463	627	1.18	5.12
Grant-added (extensive margin)	Pearl millet	318	380	1.54	10.98
Grant-added (extensive margin)	Bambara groundnut	32	73	0.66	7.93
Group means					
Grant-added (extensive margin)				1.13	8.01
Commonly-held (top 5 baseline)				2.08	22.34

Notes: "Grant-added" is the set of three crops with statistically significant positive treatment effects on the extensive margin in Figure A.25. "Commonly-held" is the five most-cultivated crops in the control group at baseline (round 0), excluding any in the grant-added set. Group means are unweighted across crops within each group.

The extensive-margin additions average skewness 1.13 and excess kurtosis 8.0, against 2.08 and 22.3 respectively for the most commonly held crops. The added crops are therefore both less right-skewed (less upside-loaded) and substantially less heavy-tailed (less exposed to rare large-loss realizations) than the established staple-and-cash-crop core of the typical baseline portfolio.

Three observations follow. First, the crop-level pattern indicates that the grant additions are not a random draw from the available menu in higher-moment space: the entries cluster among crops with thinner tails and less skewness, consistent with a fixed-cost-relaxation interpretation in which the entry-margin crops are precisely those with sufficiently stable, predictable production technology that the household can commit to a minimum-scale allocation under uncertainty. Second, despite this crop-level shift, the portfolio-weighted aggregate is unchanged because the new entries receive small land allocations (the extensive-margin entry

on Bambara groundnut, groundnut (high), and pearl millet adds at most a few percent of total acreage per household). Third, the welfare implication is that the mean-variance bound reported in the main text is not systematically biased by an omitted skewness or kurtosis dimension: in this sample and at the magnitudes documented here, the higher-moment channel adds no detectable welfare gain beyond what the variance reduction captures. A reader who places weight on the crop-level evidence might still note that the entries lower portfolio left-tail exposure (lower excess kurtosis at the crop level), which would be a welfare gain that the variance bound understates; symmetrically, the entries reduce portfolio right-skewness, which the variance bound correspondingly overstates. The two effects are of similar magnitude and offsetting sign at the household-portfolio level, which is why the formal ANCOVA finds a portfolio-aggregated null.

G.5: Robustness to Baseline Imbalance

Table B.2 reports a statistically significant baseline difference between the grant arm and the pure-control arm on portfolio standard deviation: the grant group's mean baseline Portfolio SD exceeds the control's by 9.16 USD/HA, more than twice the absolute magnitude of the headline ANCOVA treatment effect. A baseline imbalance of this size is unusual for a randomized design and merits formal investigation: while ANCOVA controls for the baseline level of the outcome and is the preferred specification under McKenzie's (2012) precision argument when within-household autocorrelation is high, the size of the imbalance relative to the treatment effect raises the question of whether the headline result is sensitive to baseline-level absorption. This appendix reports four diagnostic exercises; Table G.7 collects the first three and Table G.8 the fourth.

Randomization inference on the imbalance. The Beaman, Karlan, Thuysbaert, and Udry (2023) replication archive contains $10,000$ pre-computed permuted treatment assignments that respect the full village-level and household-level randomization structure of the original experiment, including all stratification variables. For each permuted assignment, I compute the difference in baseline Portfolio SD between the permuted-treatment and permuted-control groups, generating a sampling distribution of baseline imbalances under the maintained null that the actual randomization is correct and the observed imbalance is a chance draw. Across the full $10,000$ permutations, the absolute difference equals or exceeds the observed $|9.16|$ USD/HA in 0.0294 (two-sided) of cases. The imbalance is therefore an unusual draw rather than a typical one, but a two-sided RI p of 0.0294 is the order of magnitude of an unlucky randomization under a correct design, not of a randomization failure.

Difference-in-differences as a complement. McKenzie's (2012) ANCOVA-vs-DiD comparison favors ANCOVA on precision grounds in panels with positive within-household autocorrelation, but ANCOVA's coefficient on the baseline level of the outcome is generally less than 1, so a baseline imbalance is partially

absorbed rather than fully removed. Under a difference-in-differences specification, the change in Portfolio SD from baseline to follow-up is regressed on the grant indicator and round fixed effects, with no separate baseline-level term: $\Delta\text{PortSD}_{i,t} = \beta\,\text{Grant}_i + \cdot + \varepsilon_{i,t}$. This eliminates baseline-level dependence by construction. The DiD coefficient on Grant is -10.558 USD/HA ($p = 0.0002$), larger in absolute magnitude than the ANCOVA estimate of -4.011 ($p = 0.0286$). That gap is not evidence that ANCOVA understates the effect; it is what differencing does to a chance imbalance. When the baseline and follow-up outcomes are autocorrelated at $\rho < 1$, subtracting the entire baseline gap removes more than the imbalance transmits to the follow-up period, so the DiD estimate absorbs $(1 - \rho)$ times the baseline gap as if it were treatment. Here $\rho = 0.30$, which predicts a DiD–ANCOVA wedge of -6.37 USD/HA; the observed wedge is -6.55. The DiD column is therefore reported to show how much work the baseline adjustment is doing, not as a competing estimate of the treatment effect, and the headline ANCOVA remains the estimator the design calls for (McKenzie, 2012).

Direct controls for the imbalanced crop shares. The crop-share rows of Table B.2 attribute the PortSD imbalance to elevated baseline shares of groundnut (high-intensity), rice, and cotton in the grant arm. Adding these three crop shares as additional baseline controls in the ANCOVA specification, alongside the existing baseline-PortSD control, yields a treatment effect of -4.273 USD/HA ($p = 0.0154$), essentially indistinguishable from the headline of -4.011. The imbalance in baseline shares of the high-variance crops does not therefore propagate into the treatment-effect estimate through any residual variation that the headline specification's PortSD-baseline control fails to absorb.

Three readings follow. First, the baseline imbalance is real but not extraordinary: it is distinguishable from a typical BKTU randomization draw (RI $p = 0.0294$ for portfolio SD, $p = 0.0664$ for expected profit), though the order of magnitude is that of an unusual draw rather than a failure of randomization. Second, the headline treatment effect on portfolio SD survives every absorption strategy: -4.011 under ANCOVA and -4.273 once the three imbalanced crop shares enter directly as controls. Third, running the same three diagnostics on expected profit per hectare, the other imbalanced outcome, produces the symmetric picture: the ANCOVA estimate is -2.376 USD/HA ($p = 0.1768$) and the crop-share-controlled estimate is -2.428 ($p = 0.1477$), so the return margin is not detectably different from zero under the specification the design calls for, while the risk margin is. In both cases the DiD column moves by very close to what mean reversion of a chance imbalance predicts (-6.55 against -6.37 for portfolio SD; -4.48 against -4.12 for expected profit), which is the reason it is reported as a diagnostic rather than as an alternative estimate of either effect.

Inverse-probability reweighting. The fourth diagnostic asks what the headline estimates look like on a sample forced to be balanced. I fit a logit of grant assignment on the 55 baseline variables of Table

Table G.7: Robustness to Baseline Imbalance: Portfolio SD and Expected Profit

	ANCOVA (headline)	DiD complement	ANCOVA + crop ctrl
Panel A: Portfolio SD per hectare (baseline imbalance +9.16, RI p = 0.029)			
Grant	-4.011**	-10.558***	-4.273**
	(1.832)	(2.838)	(1.763)
N	6276	6276	6276
Panel B: Expected profit per hectare (baseline imbalance +6.68, RI p = 0.066)			
Grant	-2.376	-6.854***	-2.428
	(1.759)	(2.543)	(1.677)
N	6276	6276	6276
Baseline-of-LHS control	Yes (level)	Differenced out	Yes (level)
Imbalanced crop ctrls	No	No	Yes (gnut_h, rice, cotton)
Author controls + round factor	Yes	Yes	Yes
Community FE	No	No	No
SEs clustered at HH	Yes	Yes	Yes

Mean reversion. The DiD column is reported to show the size of the baseline adjustment, not as a competing estimate: with a chance baseline gap and baseline–follow-up autocorrelation ρ, differencing overcorrects by $(1-\rho)\times$gap. Here $\rho = 0.30$ for portfolio SD and 0.38 for expected profit, predicting DiD–ANCOVA wedges of -6.37 and -4.12; the observed wedges are -6.55 and -4.48.

Notes: Column 1 is the headline ANCOVA specification on portfolio standard deviation, estimated on the authors' preferred sample over treatment arms 4 and 5 in follow-up rounds 1–2, with the Beaman et al. (2023) controls, round indicators and standard errors clustered at the household level. Column 2 is the difference-in-differences complement, with $\Delta\text{PortSD} = \text{PortSD}_t - \text{PortSD}_0$ as the LHS and no separate baseline-level control. Column 3 adds three baseline acre-share controls for the imbalanced crops (high-intensity groundnut, rice and cotton) alongside the headline baseline-PortSD control. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B.2, with indicators for missing values (71 regressors in all; McFadden $R^2 = 0.01$, as randomization implies), form inverse-probability weights for the average treatment effect, both raw and stabilized by the marginal assignment probability, and winsorize them at the 99th percentile (32 households). Households observed in the follow-up rounds without a baseline row, which the ANCOVA carries with the missing-baseline indicator, keep a weight of one (27 household-rounds). Reweighting mechanically removes the imbalance: the standardized mean difference on baseline portfolio SD falls from 0.137 to 0.001 (0.003 stabilized), and the mean absolute standardized difference across all 55 variables from 0.039 to 0.005. Table G.8 re-estimates the headline ANCOVA under each weight for every outcome in the main and efficiency suites. The portfolio-SD effect is −4.01 USD/HA unweighted, −3.80 reweighted ($p = 0.033$), and −3.88 under stabilized weights ($p = 0.029$); across all outcomes no estimate moves by more than 0.47 of its headline standard error. I see this stability as more supporting evidence that the imbalance is not a threat to my findings.

Table G.8: Headline Treatment Effects under Inverse-Probability Reweighting

Suite	Outcome	ANCOVA (headline)	IPW	Stabilized IPW
Main	Cultivated area (HA)	0.363*** (0.114)	0.359*** (0.117)	0.360*** (0.117)
Main	Number of crops	0.333*** (0.069)	0.328*** (0.070)	0.329*** (0.070)
Main	Normalized Herfindahl index	-0.872** (0.434)	-0.753* (0.426)	-0.780* (0.428)
Main	Total profit (USD)	74.437** (32.314)	71.446** (33.596)	71.496** (33.559)
Main	Profit per hectare	4.487 (3.164)	4.468 (3.252)	4.496 (3.249)
Main	Portfolio expected return / HA	-2.376 (1.759)	-2.014 (1.756)	-2.044 (1.759)
Main	Portfolio SD / HA	-4.011** (1.832)	-3.804** (1.779)	-3.881** (1.780)
Efficiency	Sharpe ratio	0.668** (0.259)	0.597** (0.260)	0.600** (0.260)
Efficiency	Distance to crop-set frontier	-1.482*** (0.532)	-1.488*** (0.522)	-1.501*** (0.524)
Efficiency	Distance to crop-set k-frontier	-0.246 (0.471)	-0.299 (0.462)	-0.285 (0.464)
Efficiency	Sharpe at closest frontier point	0.090 (0.109)	0.096 (0.099)	0.104 (0.099)
Efficiency	Portfolio variance, diagonal	-2678.238** (1219.482)	-2681.670** (1127.668)	-2769.994** (1129.256)
Efficiency	Portfolio variance, off-diagonal	176.472 (430.805)	282.243 (435.519)	299.115 (435.812)
Efficiency	Selection Ratio	0.676 (0.946)	0.271 (0.954)	0.231 (0.955)
Efficiency	Diversification Ratio	1.155*** (0.345)	1.101*** (0.343)	1.127*** (0.345)

Notes: Grant coefficient, standard error in parentheses, from the headline ANCOVA specification of Section 6 (the Beaman et al. (2023) controls, round indicators, a baseline control for the dependent variable, standard errors clustered at the household), unweighted in the first column and weighted by inverse propensity scores in the second and third. The propensity score is a logit of grant assignment on the 55 baseline variables of Table B.2 with missing-value indicators; raw weights are $1/\hat{p}$ for grant and $1/(1-\hat{p})$ for control households, stabilized weights multiply these by the marginal assignment probabilities, and both are winsorized at the 99th percentile. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

G.6: Identification of the Fixed-Cost Mechanism against Plausible Alternatives

The size-diversification gradient documented in Section 5.6 (doubling cultivated area is associated with adding approximately one crop type) is consistent with a per-crop fixed cost, but it is also potentially consistent with three alternative mechanisms in which the same gradient would emerge without any indivisibility in crop allocation: soil heterogeneity that increases with farm size (larger farms span more soil and microclimate types, making more crops agronomically suitable); labor supply (larger households farm more land and have more family members available to supervise additional crops); and learning or experience (older, wealthier farmers who have accumulated more crop-specific knowledge cultivate larger farms and a wider menu). Because all three alternatives also predict weaker treatment-effect responses on the smallest farms, the Q1 null result reported in Figure A.27 is suggestive of the fixed-cost mechanism but does not by itself distinguish between these competing explanations.

This appendix reports two diagnostics that have differential implications under the three mechanisms. The first compares the gradient across communities with low versus high within-community heterogeneity in baseline profit per hectare; if the gradient were driven by soil heterogeneity, then communities with greater within-community productivity dispersion (more soil and microclimate variation) should display a steeper gradient. The second tests whether the gradient flattens discontinuously at small farm sizes; a binding entry constraint at small K predicts such a flattening, while the smooth alternative mechanisms (soil, labor, learning) do not. Both diagnostics use *pre-floor* cultivated area $K = \sum_c \text{PlotArea}_c \cdot \text{rev_ratio}_c$, computed

before the 0.2 HA measurement-error floor that the main cleaning process applies to plot-crop areas; this avoids contaminating either the heterogeneity proxy or the piecewise-break test with the floor's mechanical inflation of small-farm K (see the floor discussion at the end of Section 5.6). Both diagnostics are reported in Table G.9.

Table G.9: Identifying the Fixed-Cost Mechanism: Soil-Heterogeneity Split and $L_{c,min}$ Discontinuity

Panel A: Soil-heterogeneity split (community-level)					
Pooled gradient (log K)	1.138***	(0.027)			
Low-heterogeneity communities	1.158***	(0.042)			
High-heterogeneity communities	1.121***	(0.035)			
$\log K \times$ high-het interaction	-0.037	(0.055)			
Panel B: Piecewise-linear gradient with break-point candidates					
Break-point (HA)	Slope below	(SE)	Slope above	(SE)	Below − above (p)
$K < 1.50$ HA (p5, $n_{\text{below}} = 348$)	0.734	(0.109)	1.172	(0.032)	-0.438*** (0.000)
$K < 2.00$ HA (p10, $n_{\text{below}} = 674$)	0.847	(0.080)	1.162	(0.034)	-0.315*** (0.000)
$K < 3.30$ HA (p25, $n_{\text{below}} = 1928$)	0.971	(0.062)	0.995	(0.048)	-0.025 (0.734)
$K < 5.16$ HA (p50, $n_{\text{below}} = 3855$)	1.121	(0.043)	0.724	(0.069)	0.397*** (0.000)

Notes: Panel A regresses the household's baseline number of cultivated crops on $\log K$ where K is total cultivated area, with community fixed effects and community-clustered standard errors, separately for low- and high-heterogeneity communities (median split on the within-community cross-farmer SD of baseline profit per hectare, demeaned by crop). Panel B reports a piecewise-linear specification with the gradient estimated separately below and above each candidate break point; the last column reports the below-minus-above difference with the p-value of its equality test (delta method), and the stars refer to that test. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Soil-heterogeneity split (Panel A). The pre-floor pooled gradient is 1.138 (highly significant), and the pooled estimate decomposes into 1.158 in low-heterogeneity communities and 1.121 in high-heterogeneity communities. The interaction coefficient on $\log K \times \mathbf{1}$[high heterogeneity] is -0.037 (SE 0.055, $p = 0.500$), indistinguishable from zero. Under the soil-heterogeneity alternative the interaction would be reliably positive: communities with more variable agronomic conditions would absorb additional crop entries more easily as farm size grows. The data therefore reject the soil-heterogeneity story as the operative driver of the gradient: communities at the high end of the productivity-dispersion distribution exhibit no greater gradient than communities at the low end, despite presumably spanning a wider menu of soil and microclimate types.

Small-farm flattening (Panel B). The piecewise-linear specification estimates the gradient separately below and above four candidate break-points selected at the 5th, 10th, 25th, and 50th percentiles of pre-floor cultivated area. A binding entry constraint at small farm sizes predicts a flatter gradient below a low-percentile break and no discontinuity at higher percentiles; the smooth alternative mechanisms predict no break anywhere. At the 5th percentile of pre-floor K ($K < 1.50$ HA, $n = 348$), the gradient below the break is 0.734 versus 1.172 above, with the slope difference highly significant ($p = 0.000$). The same pattern holds at the 10th percentile ($K < 2.00$ HA): slope below 0.847, slope above 1.162, $p = 0.000$. At

the 25th percentile the slope difference is no longer significant ($p = 0.734$), consistent with the constraint operating only in the lower tail of the farm-size distribution. At the median ($K < 5.16$ HA), the gradient flips: it is now significantly steeper below than above ($p = 0.000$), consistent with a menu-exhaustion effect at large farm sizes where households cultivate most of the available crop set and additional area provides little extensive-margin scope. Both the small-farm flattening and the upper-tail menu exhaustion sit at the percentiles where their respective mechanisms (a binding entry constraint, a finite menu) are most plausibly operative, and neither appears at the 25th percentile.

Overall Interpretation of Table G.9 The combination of a null soil-heterogeneity interaction (Panel A) and a significant gradient flattening at small farm sizes (Panel B) is jointly more consistent with a binding entry constraint, of which a per-crop fixed cost is one example, than with the smooth soil-heterogeneity alternative. The diagnostics here do not rule out partial contributions from labor or learning, both of which can generate similar small-farm flattening patterns; what they do rule out is that the gradient as a whole is a soil-heterogeneity artifact, and what they identify is a structural break at small farm sizes that smooth alternatives by themselves would not produce. Relative to the Q1 null result reported in Section 5.6, which is consistent with all three mechanisms, this evidence tightens identification toward the entry-constraint family.

Intercropping and the plot margin. A crop added by intercropping an existing plot requires seeds but not fresh plot preparation, so the per-crop loadings reported in Section 5.6 could in principle be tracking plots rather than crops. Re-estimating with the number of plots controlled leaves the seed loading unchanged at \$4.38 per crop, with plots loading nothing, and the plough loading at \$3.48 per crop alongside a small per-plot component (\$0.93). Fertilizer's crop and plot margins are too collinear to separate (the within-farmer correlation of crop and plot counts is 0.73), so its null is stated on the unconditional crop-count margin only.

One data convention behind the size-diversification gradient reported in the main text is recorded here, because the gradient depends on it. The cleaned plot-crop area variable used in this paper applies a 0.2 HA floor, $\text{area} = \max(\text{PlotArea} \cdot \text{rev_ratio}, 0.2)$, to address measurement error in tiny plots where reported areas of 0.01–0.05 HA are likely instrument noise rather than genuine cultivation (Carletto et al., 2015). This floor inflates total cultivated area K for households with many small plots, by 6.7% on average at baseline (median 3.8%, adding 4.5% to aggregate cultivated area), so a direct estimate of the per-crop minimum-area threshold $L_{c,\min}$ from cleaned data conflates the structural threshold with the measurement-error correction. The size-diversification gradient $\hat{\beta}_{\log K} = 1.138$ reported above is estimated on pre-floor cultivated area ($K = \sum_c \text{PlotArea}_c \cdot \text{rev_ratio}_c$, summed before any floor is applied) for exactly this reason. Re-estimating the same specification on the post-floor variable yields a gradient of 1.278, roughly 12% larger; this inflation is mechanical, reflecting the floor's lift of small plot-crop areas to 0.2 HA, which raises K disproportionately

for households with many small plots, rather than a genuinely steeper relationship between farm size and cultivated cardinality. The qualitative claim is unchanged across the two conventions, but the headline magnitude is intentionally taken from the cleaner pre-floor estimate. The fixed-cost identification diagnostics reported in Appendix G.6 use pre-floor area throughout to avoid this contamination. The qualitative claim of a binding entry constraint at small farm sizes does not depend on the specific magnitude of $L_{c,\min}$: the small-farm gradient flattening reported in Appendix G.6 Panel B holds at distribution-tail percentiles regardless of how the threshold is parameterized.

G.7: Measuring the Covariance Matrix without Estimating the Entry Cost

The framework in this paper treats a per-crop entry cost as the reason households hold few crops, and it estimates a crop-level covariance matrix $\mathbf{C}$ directly. It does not estimate the entry cost itself. That asymmetry is deliberate, and it follows from what the data record rather than from a modeling preference. The two objects sit at different units of observation. Crop returns are realized outcomes measured where the crop is the unit: profit per hectare at the plot-crop-year level, from which $\mathbf{C}$ is formed as a second moment after removing farmer and year effects. Input expenditure, by contrast, is recorded at the plot level, and plots are frequently intercropped. A per-crop cost is therefore not something these data measure for every plot; it is something an allocation rule must construct on all intercropped plots. Any figure I reported for the cost of adding a crop would inherit the allocation rule rather than describe the technology.

The two objects also differ in whether a level has to be believed. $\mathbf{C}$ enters the analysis only through the portfolio variance $w'\mathbf{C}w$ and the curvature of the resulting frontier, both of which depend on how crop returns comove and not on the level at which any one crop's profit is recorded. Classical measurement error in levels largely cancels from a correlation, and the within-crop winsorization applied in construction bounds the influence of extreme plots without shifting central tendency. An entry cost is a level, and in an additive specification $E_i = Fn_i + cA_i + \varepsilon_i$ it is identified by extrapolation to $n_i = 0$, a farm that cultivates nothing and appears nowhere in the sample. Consistent with that fragility, the implied per-crop cost moves between single digits and roughly \$32 per crop across defensible treatments of controls and fixed effects (Table G.10), and it changes sign when cultivated area and cotton enter the specification without farmer fixed effects.

Two facts about costs are nonetheless well identified, and both are relative rather than absolute. First, expenditure rises with crop count: within a household, holding land and year fixed, a one percent increase in the number of crops cultivated is associated with 0.50% higher input expenditure, and 0.35% among households that do not grow cotton. Second, that increase is concentrated where a per-crop entry cost should fall and is absent where it should not: seed and plot-preparation expenditure load on the crop count

Table G.10: Implied Per-Crop Cost F under Control Escalation

	Raw	+ Area	+ Cotton	+ Village FE	+ Farmer FE
F ($ per crop)	28.40***	-104.72***	-146.65***	8.41**	31.95***
	(3.95)	(4.03)	(3.75)	(3.75)	(4.75)
N	20,644	20,644	20,644	20,644	20,593

Notes: Each column regresses winsorized household input expenditure on the crop count under progressively richer specifications; the crop-count coefficient is the implied per-crop cost F under the additive model $E_i = Fn_i + cA_i + \varepsilon_i$. The instability across columns, including sign reversals when area and cotton enter without farmer fixed effects, is the point: F is an out-of-sample extrapolation to $n = 0$ and is not credibly identified in levels. Standard errors clustered at household level in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

while fertilizer loads on cultivated area, the line-item contrast reported with magnitudes in Section 5.6 (Table G.11).

What the data do not support is the stronger claim that the entry cost is constant in the number of crops. Under the additive specification the log-log coefficient and $F\bar{n}/\bar{E}$ estimate the same quantity, namely the share of expenditure that is fixed. Estimated separately they are 50% and 17%, a factor of 2.9 apart, and the increments in expenditure across adjacent values of the crop count decline rather than repeat. I therefore use the entry cost as a modeling device that generates a cardinality constraint on the portfolio problem, discipline it against the extensive-margin behavior it predicts, and report no estimate of its magnitude. The pattern the data do deliver, a substantial cost of moving from a single crop to two that dissipates as further crops are added, speaks to why the least diversified households remain undiversified. It does not explain where an already diversified household stops, and I do not claim it does.

Table G.11: Input Expenditure by Category: Per-Crop against Per-Hectare Scaling

Category	Expected	$ per crop	SE	% of spending
Seeds	Lumpy per crop	4.28***	(0.39)	3.4
Plough / plot prep	Lumpy per crop	4.04***	(0.46)	3.5
Fertilizer	Per hectare	3.96	(3.39)	61.2
Hired labor	Per hectare	7.42***	(0.95)	7.3
Other inputs	Mixed	12.32***	(1.50)	24.5
Manure	Mixed	-0.07	(0.06)	0.1

Notes: Each row regresses household spending in one input category on the household's crop count with cultivated area held fixed and farmer and round fixed effects, so the coefficient is the change in dollars per additional crop cultivated. The *Expected* column records which way indivisibility predicts the category to scale, per crop or per hectare, and is assigned by category rather than read off the estimates. Seeds and plough or plot preparation are the two categories a per-crop entry cost predicts to be lumpy, and both are significant at the 1% level. Fertilizer, the largest category at 61% of input spending, is indistinguishable from zero on the crop-count margin ($p = 0.24$), which is the contrast the mechanism predicts. Hired labor and other inputs also load significantly on crop count, so per-crop loading is not confined to the two predicted categories; the mechanism is identified by the fertilizer contrast rather than by lumpiness being unique to seeds and ploughing. Share of spending is the category's mean share of total household input expenditure. Standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

G.8: Crop-Set Robustness: Convention and Fixed Effects

The crop-set specific efficient frontiers in Section 5.3 are constructed using the menu of crops ever cultivated by any household in each community across the four-round panel. This convention treats a crop's never-being-observed in a community as evidence of agronomic or economic infeasibility for that community. A potential concern is that never-observation across four rounds need not establish infeasibility: it could reflect high entry costs, missing knowledge, coordination failure, production requirements that local conditions cannot support, or market access. The last is the most concrete: most sales run through itinerant traders who buy raw production and resell it to warehouses and other intermediaries, and a crop that no trader visiting a community buys gives households little reason to grow it for sale. The data do not distinguish these, and the main text (Section 5.3) treats what is never grown as what is not available. The directional concern is that if grant recipients begin cultivating a crop the control arm never grew, the current convention includes that crop in the community frontier (because the panel sees it being cultivated post-baseline), which may flatter the treated household's distance-to-frontier improvement. This appendix reports two diagnostics.

The first is **count of treatment-introduced community-crop pairs.** Using the 16-crop menu of the headline analysis I compute community-crop pairs under two conventions: (i) the full-panel convention (current), where a crop is in a community's menu if any household grew it in any round; (ii) a baseline-both-arms convention, where a crop is in the menu only if it was cultivated at round 0 by either treatment or control. The full-panel convention yields 1,632 community-crop pairs across the 110 arms 4 and 5 communities; the baseline-both-arms convention yields 1,282 pairs. Most of the difference reflects pre-treatment versus post-treatment expansion that is not specific to either arm: many communities cultivate additional crops in rounds 1–3 that no household in either arm cultivated at round 0, and these account for the bulk of the 21.4% gap. Restricting to crops that the treated arm grew at any point but that no control household ever grew across all four rounds yields only 38 community-crop pairs, of which 28 were post-baseline entries. Those post-baseline entries are concentrated in six minor crops (fonio, Bambara groundnut, yam, pearl millet, sesame, sweet potato) and appear in 23 communities, 20.9% of the sample.

It is worth being precise about what these 28 pairs can and cannot do to the results, because the menu question and the moment question are separate. Each crop's μ_c and σ_c are estimated from pure-control cultivators pooled across all communities (Section 2.2), so a treated household that enters fonio is evaluated against the same fonio moments as everyone else, estimated entirely from control households and never from its own realized outcomes. A community whose control arm never grew a crop simply contributes no observations to that crop's moments; it is the community, not the crop, that drops out of the estimate. The treated household therefore adopts the crop's risk and return as measured elsewhere in the control arm, and

the only channel by which the current convention could flatter its measured efficiency is menu inclusion: whether the crop counts as feasible for that community at all. That channel is what the 28 pairs bound, and the alternative-convention treatment effect reported next is what tests it.

The second is **frontier-distance treatment effect under the alternative convention.** For each community, I rebuild the cardinality-unconstrained efficient frontier from its baseline-both-arms crop set, recompute each household-year's distance to that smaller community frontier, and re-run the headline ANCOVA on the resulting distance measure. The treatment effect on distance-to-frontier under the baseline-both-arms convention is -158.71 (SE 55.95, $p = 0.005$), even larger in absolute magnitude than the headline estimate of -142.47. The qualitative conclusion, that the grant moves households meaningfully closer to their crop-set specific efficient frontier, is preserved, and is if anything strengthened by the more restrictive crop-set convention. The mechanical reason is that the alternative convention shrinks the comparison frontier without symmetrically shrinking the household's actual portfolio, so a treatment-induced reallocation that raises Sharpe and lowers σ continues to pull the household's (σ, μ) position toward the smaller frontier.

The headline ANCOVA omits fixed effects beyond a round indicator, since the experiment was randomized within (no-loan) villages and the original Beaman et al. (2023) authors absorb village fixed effects only to handle the loan-village stratification. A diagnostic concern is that communities differ in the crop sets households actually cultivate: the 110 arms 4 and 5 communities in the analysis sample contain 18 unique crop availability patterns over the 16 analysis crops, ranging from 1 community to 39 communities per pattern. Cross-community mean-variance comparisons are therefore not perfectly apples-to-apples, motivating a fixed-effect robustness check.

Table G.12 reports Grant treatment effects on every outcome in the main and efficiency suites under three specifications: no fixed effects (the headline), the 17-dummy crop-set fixed effect, and the 110-level community fixed effect. Most outcomes are essentially invariant across the three specifications. Cultivated area, the number of crops, the Herfindahl indices, profit per hectare, portfolio standard deviation, the Sharpe ratio, the community-frontier distance measures and the variance decomposition columns all give qualitatively the same point estimates and the same sign-and-significance conclusions regardless of which fixed effect is used. Two of these, the normalized Herfindahl index and the distance to the crop-set frontier, shed one significance star under the 110-level community fixed effect while remaining significant at conventional levels; that is the precision cost of absorbing 110 dummies rather than a change in what the estimates say.

The substantive exception is portfolio expected return. Without fixed effects, the grant treatment effect on portfolio expected return per hectare is -2.38 USD/HA ($p = 0.177$, not statistically distinguishable from zero). Under the principled crop-set fixed effect, the estimate is -2.69 USD/HA ($p = 0.112$, also not significant). Under the 110-level community fixed effect the point estimate is -2.99 USD/HA ($p = 0.043$). The

headline maintained-zero hypothesis on $\Delta\mu$ is preserved under both the headline and the crop-set fixed-effect specification; the community fixed-effect specification is the only one that crosses conventional significance, and it does so by absorbing within-crop-set variation that the experimental design does not require to be absorbed.

The portfolio-variance load-bearing result is robust across all three specifications: ΔPortSD = -4.01 USD/HA without FE ($p = 0.029$), -4.23 under crop-set FE ($p = 0.019$), and -4.05 under community FE ($p = 0.013$), all significant at conventional levels.

G.9: Baseline Wealth Index: Construction, Censoring, and Ag-Asset Robustness

Why a new wealth index. The ag-asset index pre-registered for the original experiment, and built for that study's questions rather than for portfolio analysis, is severely right-censored: 31% of households pile at the ceiling value ($z = 2.258$). Within the top tercile the index takes only 3 unique values, so the top tercile is essentially a tie. Pre-grant heterogeneity in the upper half of the wealth distribution is therefore unmeasurable with that index, which is why the headline heterogeneity in Section 6.3 uses the constructed index.

Construction. The wealth index (`rich_sum` in the replication code) is an unweighted sum of nine z-scored baseline (round-0) endowment variables, drawn from the household roster and deliberately excluding choice-like variables such as cultivated area and labor days. The components are:

- *Productive durables* (per capita, extended household; linear z): plough, cart, motorbike.
- *Consumption durables* (per capita, extended household; linear z): radio, stove, television.
- *Livestock* (linear z): total animals owned.
- *Asset values* ($\log(1+x)$ then z): owned-compound value and other durable-asset value, logged because both are heavily right-skewed; without the log, the linear z compresses the poor end into a flat band.

The resulting index has no households at the ceiling and 50 unique values within the top tercile, and correlates with the pre-registered ag-asset index at $r = 0.32$. Household-round observations are sorted by the household's baseline index into three equal-sized bins ($n_{T1} = 2,085$, $n_{T2} = 2,085$, $n_{T3} = 2,085$), and the headline ANCOVA is run separately within each.

Robustness: the ag-asset index. Table G.13 reproduces the tercile heterogeneity on the censored ag-asset index. Read on its own, the ag-asset index (Table G.13) tells the tidy DARA-plus-threshold story. It reveals a monotone-declining gradient on the Sharpe ratio and the efficiency distances, and a monotone-rising gradient on the number of crops cultivated, as if the intensive envelope collapsed from poor to rich while the extensive margin strengthened. What the ag-asset index cannot show is the top-tercile detail that the

Table G.12: Crop-Set FE Robustness: Main and Efficiency Suite Outcomes

Suite	Outcome	TE (no FE)	SE	TE (Crop-set FE)	SE	TE (Comm FE)	SE
main	Cultivated area (HA)	0.363***	(0.114)	0.363***	(0.112)	0.370***	(0.106)
main	Number of crops	0.333***	(0.069)	0.330***	(0.066)	0.305***	(0.060)
main	Herfindahl index (acres)	0.000	(0.000)	0.000	(0.000)	0.000	(0.000)
main	Normalized Herfindahl index	-0.872**	(0.434)	-0.876**	(0.422)	-0.761*	(0.413)
main	Total profit (USD)	74.437**	(32.314)	70.141**	(31.959)	69.439**	(31.339)
main	Profit per hectare	4.487	(3.164)	3.794	(3.115)	3.693	(2.977)
main	Portfolio expected return / HA	-2.376	(1.759)	-2.688	(1.689)	-2.991**	(1.474)
main	Portfolio SD / HA	-4.011**	(1.832)	-4.231**	(1.804)	-4.053**	(1.629)
efficiency	Sharpe ratio	0.668**	(0.259)	0.625**	(0.243)	0.449**	(0.218)
efficiency	Distance to crop-set frontier	-1.482***	(0.532)	-1.418***	(0.515)	-1.172**	(0.480)
efficiency	Distance to crop-set k-frontier	-0.246	(0.471)	-0.242	(0.461)	-0.070	(0.427)
efficiency	Sharpe at closest frontier point	0.090	(0.109)	0.089	(0.107)	0.063	(0.101)
efficiency	Portfolio variance, diagonal	-2678.238**	(1219.482)	-2716.039**	(1210.782)	-2359.500**	(1133.442)
efficiency	Portfolio variance, off-diagonal	176.472	(430.805)	132.950	(405.087)	-25.709	(367.071)
efficiency	Selection Ratio	0.676	(0.946)	0.639	(0.910)	0.722	(0.826)
efficiency	Diversification Ratio	1.155***	(0.345)	1.137***	(0.333)	1.098***	(0.308)

Notes: Grant treatment effects on each outcome under three fixed-effect specifications. "No FE" is the headline ANCOVA with round indicator and authors' control suite. "Crop-set FE" is a 17-level categorical fixed effect identifying communities by their unique crop availability pattern across the 16 analysis crops. "Comm FE" is a 110-level community fixed effect. Standard errors clustered at the household level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

constructed index exposes. Because 31% of households pile at the index ceiling and the top tercile takes only 3 unique values, the censoring collapses T3 into a near-tie and masks the portfolio-SD reduction the richest households achieve through crop entry; the both-tails pattern in Table B.15 is therefore invisible under the pre-registered index. The two indices nonetheless agree on the robust features: the extensive-margin crop response rises with wealth under either measure, and the poorest tercile posts the largest within-cardinality efficiency gains in both. The disagreement is confined to the upper tail, which is precisely where the ag-asset index has no resolving power; Figure G.2 repeats the exercise by ag-asset decile.

Table G.13: Robustness: Heterogeneous Treatment Effects by Baseline Ag-Asset Wealth Tercile

Outcome	T1 (poorest)	T2 (middle)	T3 (richest)
Number of crops	+0.2	+0.4***	+0.4***
Expected return / HA	−1.3	−0.8	−4.4
Portfolio SD / HA	−4.3	−3.6	−3.6
Sharpe ratio (×100)	+1.1***	+1.0**	0.0
Distance to community frontier	−2.1***	−1.7**	−0.6
Distance to community k-frontier	−1.7**	−0.1	+1.0
Selection Ratio (self-insurance)	+0.7	+0.3	+1.1
Diversification Ratio	+0.5	+1.4**	+1.4**
n (HH-rounds)	2,086	2,085	2,082

Notes: Terciles formed on baseline (round-0) ag-assets. The censored ag-asset index cannot resolve upper-tail heterogeneity; see the construction of the wealth index above and the headline table (Table B.15). * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

G.10: The Frontier-Projection Welfare Bound: Derivation, Computation, and CRRA Interpretation

This appendix derives the certainty-equivalent welfare bound reported in Section 7 and gives the full computation. The method adapts the frontier-projection logic of Calvet et al. (2007) to *bound*, rather than point-identify, each household's risk aversion.

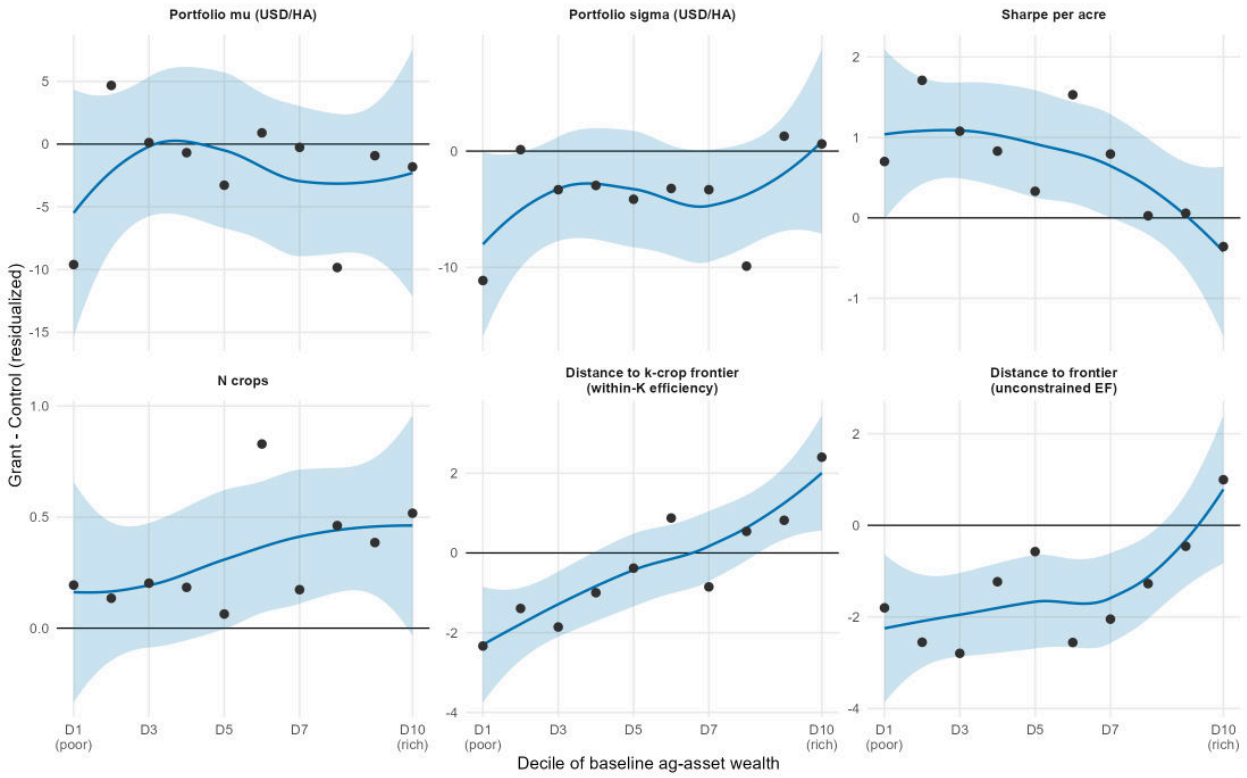
Setup. Under mean-variance utility $U(\mu, \sigma) = \mu - \frac{\gamma}{2}\sigma^2$, the certainty equivalent of a portfolio with mean μ and variance σ^2 equals U itself (Levy and Markowitz, 1979), so any treatment-induced change delivers a welfare change $\Delta U = \Delta\mu - \frac{\gamma}{2}\Delta\sigma^2$ measured in the same dollars-per-hectare units as expected profit.

Step 1: Frontier projections. For a baseline household at an interior point (σ_0, μ_0), define two projections onto its crop-set specific efficient frontier:

$$\text{West: } (\sigma_W, \mu_0),\ \sigma_W = \min\{\sigma : (\sigma, \mu_0) \in \text{frontier}\} \qquad \text{North: } (\sigma_0, \mu_N),\ \mu_N = \max\{\mu : (\sigma_0, \mu) \in \text{frontier}\}.$$

The West projection is the most-risk-averse reading of where the household was trying to go (variance reduction at fixed return); the North projection is the most-risk-tolerant (higher return at fixed risk).

Figure G.2: Robustness: Treatment Effect by Baseline Ag-Asset Decile



Notes: Each point is the mean within-decile grant–control gap on the residualized outcome (residualized on baseline outcome, missing-baseline indicator, controls, and round FE) across baseline ag-asset deciles. Loess smooth with 95% CI.

Step 2: Implied γ from tangency. At an optimum on the frontier the mean-variance first-order condition gives $\partial\mu/\partial\sigma|_{\text{frontier}} = \gamma\sigma$, so $\gamma = (\partial\mu/\partial\sigma)|_{\text{frontier point}}/\sigma$. Evaluated at each projection this returns a farmer-level pair $(\gamma_{\text{north}}, \gamma_{\text{west}})$ bracketing the household's true γ. Across the $N = 3{,}180$ baseline households with a defined projection, the cross-farmer means are $\gamma_{\text{north}} \approx 0.00285$ and $\gamma_{\text{west}} \approx 0.00410$.

Step 3: Welfare gain from the treatment. Applying the headline ATEs to the certainty-equivalent expression, $\Delta U = \Delta\mu - \frac{\gamma}{2}\Delta\sigma^2$, where $\Delta\sigma^2$ is the ATE on the *variance* (not $(\Delta\sigma)^2$). By the delta method $\Delta\sigma^2 \approx 2\bar{\sigma}\,\Delta\sigma = 2 \times 316 \times (-4.01) \approx -2{,}532$. I use the direct regression on σ^2 throughout, which gives $-2{,}445$; the two agree to within four percent, as the second-order approximation should.

Step 4: Plug in. With the point estimate $\Delta\mu = -2.38$ (statistically insignificant) and $\Delta\sigma^2 = -2{,}445$,

$$\Delta U = -2.38 - \tfrac{\gamma}{2}(-2{,}445) = -2.38 + 1{,}222\,\gamma,$$

giving $\Delta U = -2.38 + 1{,}222 \times 0.00285 \approx \$1.11$/HA at the risk-tolerant (North) bound and $-2.38 + 1{,}222 \times 0.00410 \approx \$2.64$/HA at the risk-averse (West) bound. Using the negative point estimate rather than imposing

$\Delta\mu = 0$ trims roughly \$2.38/HA off the gain at each bound and is the conservative choice.

Step 5: Scale to grant value. Scaling per-hectare welfare by the 7.15 HA the average control household cultivates over the two estimation rounds, the same sample and rounds that produce the ATEs it multiplies, and dividing by the grant,

$$\frac{\Delta U \times 7.15}{153} = \begin{cases} 5.2\% & \text{risk-tolerant (lower) bound} \\ 12\% & \text{risk-averse (upper) bound,} \end{cases} \qquad \frac{\Delta U \times 7.15}{31} = \begin{cases} 26\% & \text{lower bound} \\ 61\% & \text{upper bound,} \end{cases}$$

where the second expression normalizes by the \$31 of incremental agricultural spending the grant induces rather than the full grant. The headline range is thus **5.2–12% of grant value** (26–61% per marginal ag-invested dollar). Imposing $\Delta\mu = 0$ rather than the point estimate leaves the γ bounds and $\Delta\sigma^2$ unchanged but sets the $\Delta\mu$ term to zero, raising the range to 16–23% of grant value (80–116% per ag-invested dollar); the reported figures are therefore conservative in the $\Delta\mu$ dimension.

Identification caveat. Under binding constraints the observed (σ_0, μ_0) reflects both risk preferences and constraint severity, so the projected γ is a composite of the two; the bound is an as-if-optimizing welfare accounting rather than a structural preference estimate. Appendix G.10 reports a complementary bound that replaces the projection with an external CRRA range.

The headline certainty-equivalent welfare bound of 5.2–12% of grant value (Section 7) is computed by projecting each baseline household's (σ_h, μ_h) onto its crop-set specific efficient frontier in two directions, recovering an implied risk-aversion parameter at each projected frontier point, and applying the ATE under each household's projected γ. A potential identification concern with this approach is that the household's observed inside-the-frontier position reflects both unobserved risk preferences and binding constraints; under the latter component, the projection identifies a composite of preferences and constraint severity rather than γ alone. A constraint-bound household with high true γ would project to a different frontier point than it would actually choose if unconstrained, and the implied γ from that projection is therefore an inconsistent estimate of the household's preference parameter.

This appendix places that projection-implied risk aversion on the conventional CRRA scale, both to interpret it and to report, for reference, the welfare gain that imposing standard CRRA values would instead give. I take the conventional development-literature range of CRRA risk aversion, $\gamma_{\text{CRRA}} \in \{1, 2, 3, 4, 5\}$, convert to the mean-variance equivalent via the Arrow-Pratt approximation, and apply the headline ATE on σ_p^2 ($\Delta\sigma^2 = -2446$ USD^2/HA^2, $\Delta\mu = 0$ maintained) directly. The units must be carried consistently. Absolute risk aversion over total profit is $\gamma_{\text{CRRA}}/\bar{c}$ with $\bar{c} = 2267$ USD/year, but the variance change here is measured per hectare, so the coefficient multiplying it is that quantity scaled by mean cultivated area,

which reduces to $\gamma_{\text{MV}} = \gamma_{\text{CRRA}}/\bar{\mu}_{\text{HA}}$, relative risk aversion over mean profit per hectare:

$$\Delta U_h = \Delta\mu - \tfrac{\gamma_{\text{MV}}}{2}\Delta\sigma_h^2.$$

Multiplying by mean farm size and dividing by the \$153 grant value yields the welfare gain as a fraction of the grant. The remainder of this section reports a complementary bound that replaces the projection with an external CRRA range.

Table G.14: Welfare Bound under Literature-Informed CRRA Risk Aversion

γ_{CRRA}	γ_{MV}	ΔU/HA (USD)	$\Delta U \cdot \bar{\text{HA}}$ (USD)	% of grant value
1	0.00311	3.80	26.78	17.5%
2	0.00622	7.60	53.57	35.0%
3	0.00933	11.40	80.35	52.5%
4	0.01243	15.20	107.13	70.0%
5	0.01554	19.01	133.92	87.5%
$\Delta\sigma = -4.011$, $\Delta\sigma^2 = -2446$, $\Delta\mu = 0$ (maintained); $\bar{c} = \mu_{\bar{\text{HA}}} \cdot \bar{\text{HA}} = 2267$ USD; grant value = 153 USD				

Notes: CRRA risk aversion converted to mean-variance via $\gamma_{\text{MV}} = \gamma_{\text{CRRA}}/\bar{\mu}_{\text{HA}}$, relative risk aversion over mean profit per hectare (derivation in the text above). Welfare gain ΔU/HA computed as $-\gamma_{\text{MV}}\Delta\sigma^2/2$ with $\Delta\sigma^2$ from the headline ATE specification (Section 6) and $\Delta\mu = 0$ maintained. Total welfare gain multiplies by mean farm size; % of grant divides by \$153. The conventional development-literature CRRA range is roughly $[1, 5]$ (Chetty, 2006); Schechter (2007) estimates median $\gamma_{\text{CRRA}} = 3.5$ for rural Paraguayan farmers, and Liu (2013) estimates median $\gamma_{\text{CRRA}} = 2.4$ for Chinese cotton farmers.

The literature-bounded welfare gain ranges from 17.5% of grant value at $\gamma_{\text{CRRA}} = 1$ to 87.5% at $\gamma_{\text{CRRA}} = 5$, with a median-of-the-development-literature estimate around 52.5% at $\gamma_{\text{CRRA}} = 3$.

Why the two estimates diverge. The divergence between the projection-based 5.2–12% range (Section 7) and the literature-bounded 17.5–87.5% range can be diagnosed by converting the projection-implied γ values back to the CRRA scale. The projection reads risk aversion off the per-hectare frontier, so it is already expressed per hectare, and recovering the CRRA scale means multiplying by mean profit per hectare rather than by total profit. The risk-tolerant North-projection bound implies a median CRRA risk aversion of $\gamma_{\text{CRRA}} = 0.9$ (interquartile range $[0.7, 1.2]$), at the lower end of the literature range. The risk-averse West-projection bound implies a median of $\gamma_{\text{CRRA}} = 1.4$ (interquartile range $[1.0, 1.6]$), also inside it.

Evaluated on the same terms, the two approaches agree. Imposing $\Delta\mu = 0$ as the literature calibration does, the projection bound is 16–23% of grant value against 17.5% at $\gamma_{\text{CRRA}} = 1$, and the projection independently recovers a median coefficient of 0.9 to 1.4. The two are reading the same households at nearly the same preference parameter, which is the consistency check the diagnostic exists to perform. The headline projection range of 5.2–12% is lower than either only because it applies the $\Delta\mu$ point estimate rather than setting it to zero, and that estimate is negative and statistically insignificant.

The higher rows of Table G.14 describe households more risk-averse than these appear to be: 52.5% at

$\gamma_{\text{CRRA}} = 3$ and 87.5% at $\gamma_{\text{CRRA}} = 5$. They are reported for reference rather than as competing estimates, since the projection places this population near the bottom of the conventional range. One caveat applies to every row: applying a single representative γ to the population-average $\Delta\sigma^2$ ignores heterogeneity in both, and the sign of the resulting bias depends on how γ_h and $\Delta\sigma_h^2$ covary across households, which this exercise does not measure.

The projection-based 5.2–12% remains the paper's welfare number; the literature-informed figures in Table G.14 are a reference point, not a co-equal bound.

G.11: Per-Hectare Risk Intensity versus Total-Farm Risk

The risk outcome used throughout the paper, $\sigma_p = \sqrt{\vec{w}'\mathbf{C}\vec{w}}$, is an intensity in dollars per hectare. The grant also raised cultivated area, so a household's total exposure need not have moved in the same direction as its risk per hectare. Within the framework the two are linked: total farm profit is $K\vec{w}'\vec{R}$ with variance $K^2\vec{w}'\mathbf{C}\vec{w}$, so total risk is $K\sigma_p$ and total expected profit is $K\mu_p$, where K is cultivated area. This section reports the headline specification run on those totals. Each number comes from estimating Equation 7 on the constructed total rather than from multiplying two separate coefficients, which would not recover the treatment effect on a product.

Table G.15 collects the result, and it splits in two. Risk intensity falls and scale rises: portfolio standard deviation drops 1.3% per hectare while area rises 5.1%. Total-farm risk therefore rises, by \$65.91 on a control mean of \$2,260 (2.9%), though only marginally, with $p = 0.065$; in variance rather than standard-deviation units the increase is 2.5% and is not distinguishable from zero at all ($p = 0.54$). Against that, total expected profit rises 3.7% (\$80.35, $p = 0.018$) and realized total profit rises 6.8% ($p = 0.021$). Return rises faster than risk, and it rises with more confidence.

The composition is therefore what the paper measures and the scale is what the grant additionally buys, and the two do not conflict. A household that holds its portfolio's risk intensity down while expanding onto more land takes on more total exposure by construction; the question is the price. Evaluating the total position in mean-variance terms, the treatment is welfare-improving for any household whose absolute risk aversion over total profit falls below $2\Delta\mu/\Delta\sigma^2$. Converting through the Arrow-Pratt scaling used in Appendix G.10, that break-even sits at a relative risk aversion of 0.97 when mean realized profit is the consumption scale and 1.90 when mean expected profit is. The projection exercise of Appendix G.10 recovers a coefficient near one, so on the paper's own estimate of risk preferences the household is close to indifferent about the scale expansion and strictly better off below that value. Nothing in the headline result depends on this comparison, since every treatment effect in the paper prices two arms with the same matrix and the

same normalization. What the section rules out is the reading that a per-hectare risk reduction was bought by quietly raising the household's total income risk on unfavorable terms.

Table G.15: Per-Hectare and Total-Farm Outcomes under the Headline Specification

	Effect	SE	p	Control mean	% of control
Per hectare					
Portfolio SD (USD/HA)	-4.01**	(1.83)	0.029	315.68	-1.27
Expected profit/HA (USD/HA)	-2.38	(1.76)	0.177	299.79	-0.79
Scale					
Cultivated area (HA)	0.36***	(0.11)	0.001	7.15	5.08
Total farm					
Total-farm risk $K\sigma_p$ (USD)	65.91*	(35.67)	0.065	2,260	2.92
Total-farm variance $K^2\sigma_p^2$ (USD^2)	182,798	(299,904)	0.542	7,362,971	2.48
Total expected profit $K\mu_p$ (USD)	80.35**	(34.07)	0.018	2,165	3.71
Realised total profit (USD)	74.44**	(32.31)	0.021	1,098	6.78

Notes: Each row is a separate regression under the headline ANCOVA specification of Section 6: the Beaman et al. (2023) control suite, round indicators, a baseline control for the dependent variable and an indicator for its absence, and standard errors clustered at the household. Sample is the two grant arms over the two estimation rounds, 6,276 household-years. Total-farm quantities are constructed per household-year as $K\sigma_p$, $K^2\sigma_p^2$ and $K\mu_p$ before estimation, so the reported effects are treatment effects on those totals rather than products of the per-hectare and area coefficients. Percent of control is the effect divided by the control mean in the same column. * p<0.10, ** p<0.05, *** p<0.01.

G.12: Dominated Crops and Monocropping Frequency

Section 5.1 notes that dominated crops, which another crop beats on both return and risk, still carry diversification value through their correlations. This section reports the test that observation implies, moved here from the main text. Several crop assets are dominated, meaning another available crop offers at least as high an expected return at no greater risk, with at least one of the two comparisons strict. Excluding them from a portfolio may seem rational at first glance, but Table B.7 shows these dominated crops still provide value: their returns are imperfectly correlated with those of other crops, offering cultivators a diversification opportunity. A testable implication is that dominated crops, valued primarily through their off-diagonal correlations rather than their own risk-return position, should rarely be cultivated as the sole crop in a household's portfolio. Empirically, household-year-level monocropping frequency does rise with each crop's Sharpe ratio, though the relationship is not statistically significant ($\hat{\beta} = 0.0066$ per Sharpe unit, $p = 0.120$, $R^2 = 0.16$ across 16 crops); the mean monocropping rate is 0.09% for the 4 dominated crops and 0.23% for the 12 undominated crops. Monocropping is rare across the board in this sample (mean 0.19%, max 0.89% for high-input maize), so the test has little power to detect a difference of this size. The direction is consistent with the diversification-value interpretation, but the gap is not statistically distinguishable from zero and I read it as suggestive rather than as strong evidence.

G.13: Multivariate Correlates of Portfolio Efficiency and the Realized-Risk Version

Section 5.5 reports the bivariate associations between baseline endowments and portfolio efficiency and their multivariate summary. This section gives the multivariate detail and the realized-σ counterpart of the risk-target exercise, both moved from the main text. Household assets, a stock measure, appears in these bivariate associations to play a smaller role than the flow and capital measures (food security, plough, fertilizer, livestock, savings) in predicting how close a household sits to its frontier. The bivariate ranking does not separate direct effects from shared variance among predictors: plough use, fertilizer use, livestock value, and farm size are themselves correlated, so a multivariate specification entering the predictors jointly would be needed to apportion predictive power. Entering the predictors jointly alters the bivariate picture meaningfully (Table B.12). Plough use, household size, farm size, and financial savings retain significance with positive signs. Three predictors that look meaningful in bivariate do not survive conditioning: fertilizer use, livestock value and food insecurity all fall to insignificance. Fertilizer use is the clearest case, strongly associated with efficiency on its own but adding little once plough use is held fixed, since the two travel together. What survives describes a household's capacity to act, through traction, land, labor, and cash on hand, rather than its accumulated stock of goods, which is the distinction the constraint reading predicts should matter. The exception is household assets, which emerges in the multivariate as a significant positive predictor ($\hat{\beta} = 0.071$, $p = 0.000$) despite a null bivariate association, a Simpson's-paradox pattern in which household assets predicts efficiency *conditional on* the correlated capital and labor measures but not in raw bivariate terms. However, the endogeneity of these covariates (plough ownership, wealth, and shock exposure are all jointly determined with agricultural decisions) limits the extent to which these associations can be given any causal interpretation.

The realized-σ version of this exercise is reported as Figure A.30. Its coefficients carry the same signs but are attenuated toward zero, because the endowments that raise a household's target also pull it closer to the frontier, and the two movements partly cancel in the realized value. Farm size is again the clearest case, entering realized risk at 0.013 ($p = 0.392$) despite moving the constraint margin substantially.

Appendix H: Long-Run Treatment Effects (2017 Follow-Up)

The main results restrict to the two post-treatment rounds immediately following the grant (2011 and 2012). The survey includes a long-run follow-up conducted in 2017, six years after the cash transfer was distributed. Table H.1 replicates the main and efficiency specifications (Panels A and B) on the 2017 data alone, testing whether the portfolio effects persist beyond the short run. The Grant coefficients on portfolio standard deviation, expected return, and the efficiency measures are statistically indistinguishable from zero throughout, indicating that the portfolio effects observed in the short run (2011–2012) do not persist six years after a one-time transfer. This is consistent with the grant providing a temporary relaxation of liquidity constraints that enabled households to restructure their portfolios in the immediate post-treatment seasons, but not generating a permanent shift in portfolio composition. It is also the pattern the fixed-cost interpretation predicts: a one-time transfer clears entry costs once, and portfolios that re-equilibrate to the household's enduring resources afterward carry no permanent imprint of the transfer itself.

The ANCOVA specification is identical to Equation 7, with the 2010 baseline value of each outcome as a control and the authors' preferred covariate suite. Because this regression uses the 2017 wave alone, no round fixed effect is included. Standard errors are again clustered at the household level.

Horizon of the welfare bound. These estimates also fix the horizon of the certainty-equivalent bound in Appendix G.10, which is a per-season flow rather than a stock. The bound applies the average treatment effect to one season's portfolio and scales it by one season's cultivated area, so the 5.2–12% of grant value it reports is what a single season delivers, set against a transfer the household receives once. The effect is estimated over the two post-treatment seasons, and the results in this appendix show it has gone by 2017. Cumulating over those two seasons and neglecting discounting across a single year puts the total at roughly 10–25% of grant value, after which the bound stops accruing. The welfare gain is therefore a temporary flow across the two seasons the grant reshaped, not a permanent change in the household's risk position, and the same one-time-transfer logic that makes the fixed-cost reading attractive is what makes the flow stop.

Table H.1: Long-Run Impacts of Cash Transfer (2017 Follow-Up)

Panel A: Measured and Estimated Outcomes

	(1) Hectares	(2) N Crops	(3) Norm HHI	(4) Total Expenditure	(5) Total Revenue	(6) Profit	(7) Profit/HA	(8) E[Profit]/HA	(9) Portfolio SD
Grant	0.14	0.13	-0.61	362.41	712.12	155.32	2.75	0.28	-0.04
	(0.17)	(0.08)	(0.82)	(294.51)	(437.87)	(130.52)	(12.99)	(2.86)	(3.13)
N	3,055	3,055	3,055	3,055	3,055	3,055	3,055	3,055	3,055
Control Mean	7.52	4.69	21.04	1989.48	5804.06	3467.73	427.23	368.99	390.88
Treatment Mean	7.58	4.82	20.38	2340.08	6432.84	3586.59	430.03	370.63	392.07

Panel B: Portfolio Efficiency

	(1) Sharpe Ratio	(2) Origin Dist	(3) Euclidean Dist, Crop-Set EF	(4) Sharpe Gap at Closest Point	(5) Selection Ratio	(6) Diversification Ratio
Grant	0.34	0.16	-0.63	-0.01	-0.04	0.34
	(0.39)	(4.09)	(0.93)	(0.30)	(1.23)	(0.54)
N	3,055	3,055	3,055	3,055	3,055	3,055
Control Mean	94.60	538.26	48.54	114.08	152.67	132.14
Treatment Mean	94.89	540.24	48.06	113.97	152.04	132.44

Notes: ANCOVA specification regressing each outcome on a treatment indicator (grant receipt), the baseline value of the outcome, an indicator for missing baseline, and the authors' preferred control suite. No community or round fixed effects. Sample restricted to grant villages (arm 4) versus pure control villages (arm 5), 2017 follow-up round only, balanced panel households. All monetary values in 2017 USD. Panel B: Sharpe Ratio is the ratio of expected return to portfolio standard deviation, scaled by 100. Origin Distance is the Euclidean distance from the origin in mean-variance space. Euclidean Dist to Crop-Set EF is the distance to the closest point on the crop-set specific efficient frontier, and Sharpe Gap at Closest Point compares Sharpe ratios at that projection. Selection Ratio is the standard deviation of the community's highest-return crop divided by the household's linear-weighted average crop standard deviation, matching the main tables. Diversification Ratio compares the weighted-average crop SD to the actual portfolio SD. Standard errors clustered at household level in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

Appendix I: Limitations of the Financial Asset Analogy

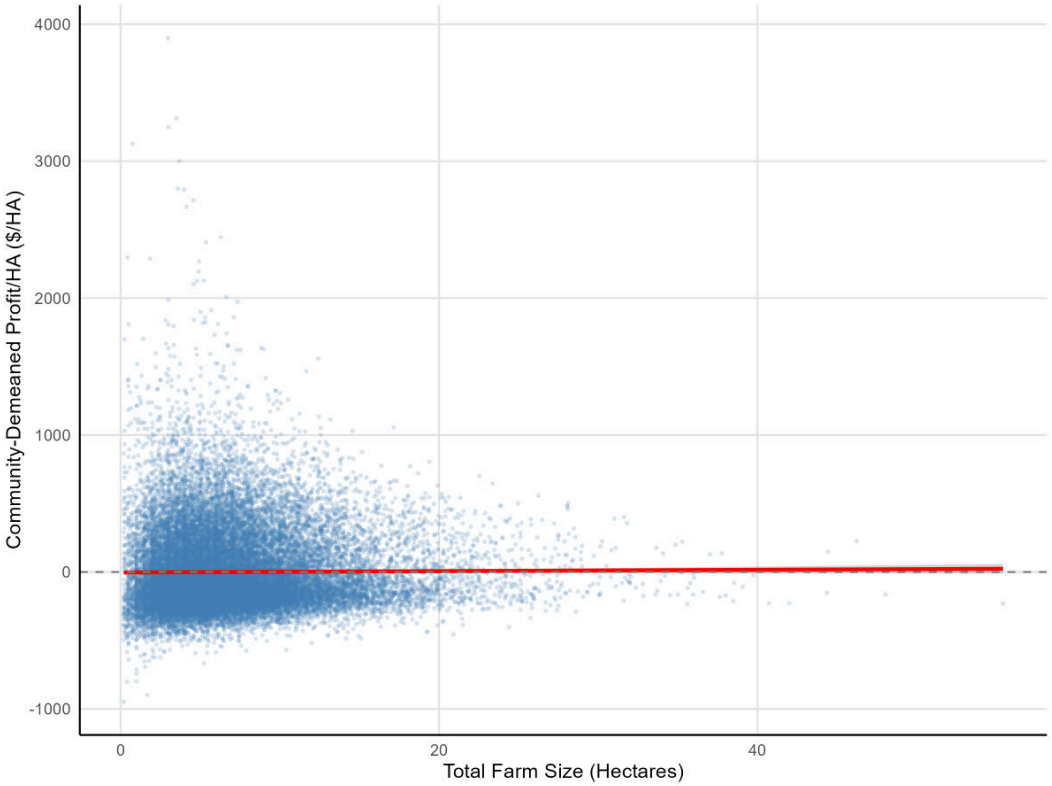
I.1: Economies of Scale

In developed countries, the crop an individual cultivates becomes an identity: the way their farm is organized, the capital they invest into and in a sense, the whole production function they obey is specialized to a particular crop like corn or soy beans. The reason for this specialization is to take advantage of economies of scale, which are well documented in agriculture (Hallam, 1991). The average size of a farm in the United States was 180.5 hectares in 2022 (National Agricultural Statistics Service, 2023). This is wildly different from what I find in my sample from Mali, where an entire community of about 29 households grow only 204 hectares on average. This tremendous difference in scale displays an ability for developed country farms to automate and mechanize processes of seeding, fertilizer application and harvesting to their selected crop of choice, optimizing at each step in production.

This benefit from economies of scale likely outweighs the diversification benefits of multicropping explored in this project. Holding weather and price risks constant, if farmers with access to economies of scale were able to earn enough profits in a single season to cover the losses of a separate shock season, the ability to self-insure across time becomes an option. Instead of needing to guarantee enough income to reach a consumption threshold each and every year regardless of the realized weather and price states, farmers could then lean on savings from prior successful seasons to smooth their consumption. This is a very different approach to agricultural production than the one implemented by smallholder farmers, who do not have the ability to finance their consumption fully across time. At what farm size exactly the consumption smoothing benefit of scale overcomes the diversification benefit to multicropping explored in this paper is an empirical question that my data unfortunately cannot answer. It remains an exciting avenue for future work.

Empirically, farmers in my sample do not appear to benefit from economies of scale: profit per hectare does not increase with farm size, as shown in Figure I.1. If economies of scale were present, I would need to account for the increased returns to holding an additional hectare of an asset in my estimation of the μ and σ crop moments. This could have been accomplished by parameterizing the expected return for each crop to the number of hectares a farmer was holding. This would have required estimating the scale benefit for each crop, dictating its increasing returns. Alternatively, I could have binned farmers by farm size and re-calculated expected return per hectare values across crops within farm-size-bin. That approach to the problem would require much more data than I have available in this paper. For practitioners working in settings where there are returns to scale present, these are two tractable methods that they can employ to account for the scale effects when applying MPT to agricultural production.

Figure I.1: No Evidence of Economies of Scale: Profit Per Hectare vs. Farm Size



Notes: Each point represents a household-season observation. Profit per hectare is community-demeaned. The red line shows the linear fit, which is effectively flat, indicating no relationship between farm size and profit per hectare in this sample.

I.2: Asset Return Complementarities through Intercropping

Unlike financial assets, that can be held in any combination without impacting each others' returns, agricultural crops planted in specific combinations have strongly documented interactions that can improve yields (Li et al., 2023). These interactions come in different forms with a classic example being that tall, leafy crops like maize will shade and steal sunlight from shorter crops with small footprints. A more complicated example is how legumes' roots hold nitrogen in the soil, which acts as a natural fertilizer for more marketable crops than legumes like maize (Abdul Rahman et al., 2021). This means that the expected return per hectare of a plot intercropped with Crop A & Crop B will not necessarily be the average return of two sole-cropped plots with the same assets. These interactions affect not only average yield levels but also yield stability across years and locations: in a meta-analysis of 33 published intercropping experiments, Raseduzzaman and Jensen (2017) find that cereal-grain legume systems deliver coefficients of variation in yield substantially lower than the corresponding sole crops, with the stability advantage holding across tropical, subtropical, and temperate climatic zones. Understanding these interactions is a first order concern for the

field of agronomy. Researchers can handle this in their applications of MPT to agricultural production by treating commonly intercropped combinations as separate assets in their own right, with moments estimated from plots planted in given combinations rather than composed from the sole-cropped constituents. Recent experimental evidence supports measuring the bundle rather than composing it: in three-year maize/pea and maize/rapeseed trials, Wu et al. (2023) find that the temporal stability of the intercropped system depends on the intensity of complementarity and competition between the two species, so its variance is not recoverable from the constituents' moments, and that intercropping does not raise stability over monoculture on average, so the sign of the interaction cannot be assumed either. I examined commonly intercropped combinations and found that far too few repeat observations of cultivation were present to enable solid estimation of the variance-properties of joint crop assets.

While this simple reclassification is a straightforward way to account for intercropping complementarities, there is unfortunately not a simple method to account for the cost-savings that come through technological complementarities across crops. These come about due to the potential for crops to prefer similar chemical inputs or have complementary labor schedules, which allows households to hire fewer additional workers. While the intercropping benefit I explain above only manifests within a plot (which makes it easier to observe, measure and account for), these cost-saving technological complementarities exist at the household level and in turn cannot be simply accounted for in my estimations of crop returns. Adjusting crop profits for these technological complementarities would require a sample of farmers who both monocropped and multicropped a given combination in different years, along with a matrix of interaction terms in the estimation of portfolio return similar to the portfolio variance estimation. I hope that one day we have the data to explore such high dimensional objects in this line of research, but I fear that there is no tractable way given presently available data to properly account for these technological complementarities.

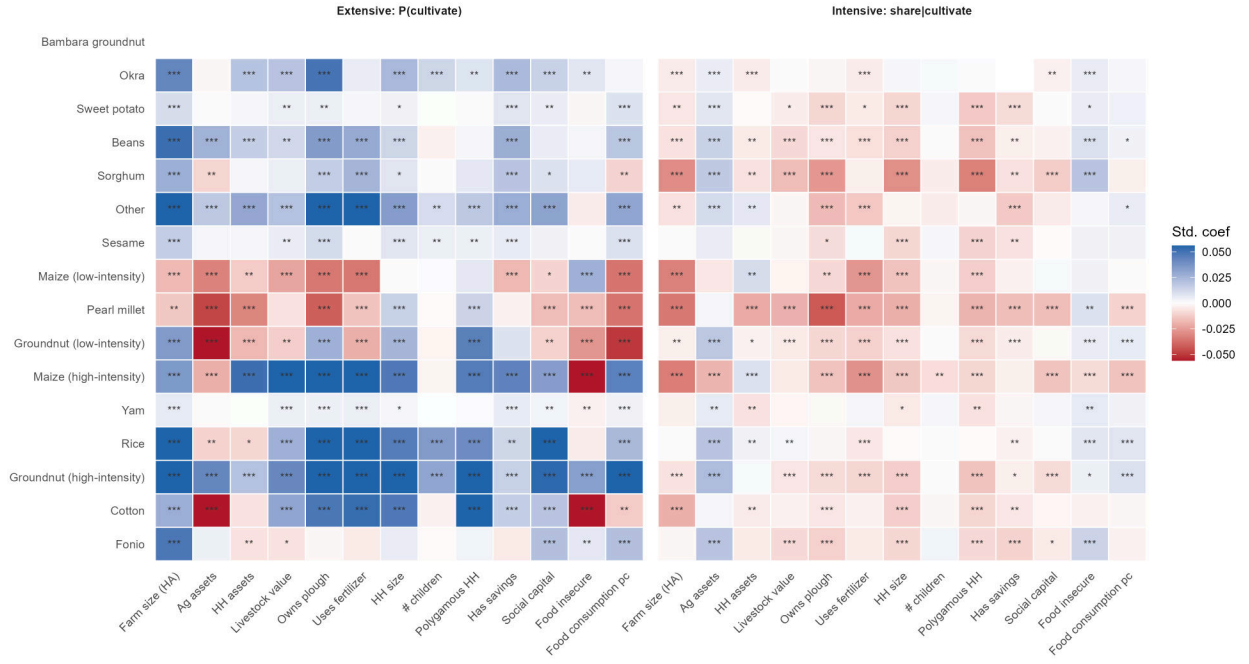
Appendix J: Cultivator Selection and Moment Robustness

The crop-level moments $(\vec{R}, \mathbf{C})$ used throughout the paper are estimated from the observed cultivations of each crop. I do not have experimental variation in crop cultivation that would inherently let me estimate the foundational $(\vec{R}, \mathbf{C})$ properties free of selection concerns. If households with different productivity levels systematically sort into different crops, the estimated moments could, in principle, reflect the ability of the cultivator composition rather than properties of the crops themselves. Section 2.4 addresses this concern at a high level; this appendix presents the supporting evidence. Three exhibits, each targeting a distinct version of this selection critique, are reported below.

J.1: Selection on observable household characteristics

Figure J.1 reports the coefficient from a bivariate OLS regression of an indicator for cultivating each crop (extensive margin) and the cultivated share conditional on positive cultivation (intensive margin) on each of 13 baseline household characteristics, run separately for each of the sixteen crop types. Crops are ordered top-to-bottom by estimated σ_c (lowest at top). Stars indicate significance at the 10%, 5%, and 1% level; cell color shows the sign and magnitude of the standardized coefficient.

Figure J.1: Household-characteristic associations with cultivation (baseline)



Notes: Each cell reports the standardized coefficient from a bivariate OLS regression of the indicated cultivation outcome (column block) on the indicated household characteristic (column), run on the baseline (2010) sample. Rows are crops ordered by estimated σ_c. Red cells = household characteristic associated with *lower* cultivation; blue = higher. Stars: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Coefficients are clipped at the 95th percentile of $|\hat{\beta}|$ for visibility. Per-crop coefficient plots for every crop are included in the replication materials.

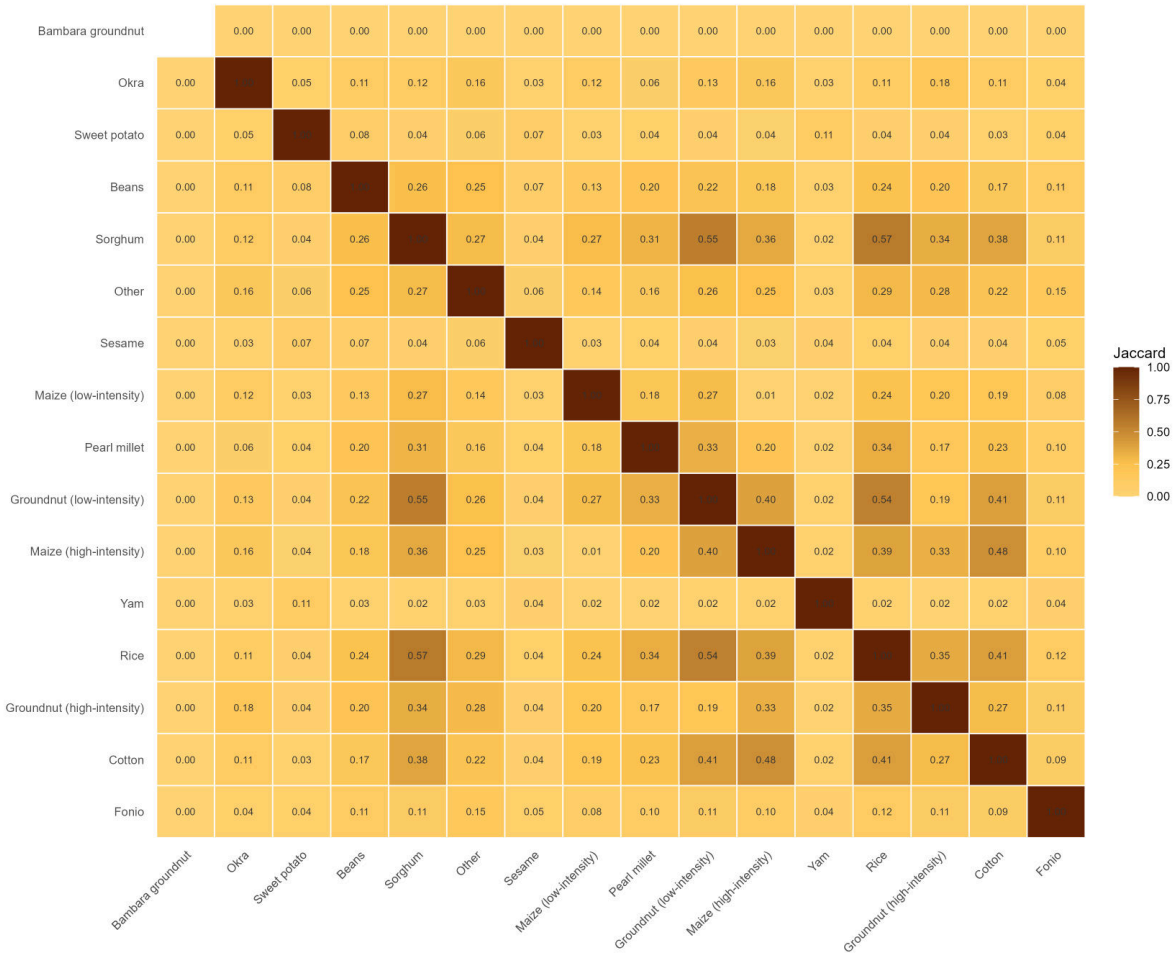
Selection on observables is real for every crop with a full set of estimates: across the 15 such crops, between 42% and 88% of the household-characteristic coefficients are significant at the 5% level depending on the crop. But the strength of the association is only weakly correlated with crop σ_c (Spearman correlation between σ_c-rank and number of significant coefficients per crop is 0.17). Where strong associations exist, they are concentrated on agronomic-infrastructure variables (rice with farm size and fertilizer use; cotton with plough ownership and household labor) rather than on wealth indicators that would proxy a risk-management capacity story.

J.2: Cultivation overlap and capability spread within crops

If selection into crop types based on cultivators' capability was severe enough to bias the moments, I would expect crops to partition into capability-based clusters This would be visible as a block-diagonal structure in a pairwise cultivation-overlap matrix, with the capable-farmer cluster sharing high cross-cultivation and the less-capable cluster sharing high cross-cultivation, with minimal cross-cluster overlap. Figure J.2 shows the pairwise Jaccard cultivation-overlap matrix. The structure does not exhibit a block partition: the median

pair of crops shares 10% cultivator overlap, and crops do not separate into "capable" and "less-capable" menus.

Figure J.2: Pairwise Jaccard cultivation-overlap matrix

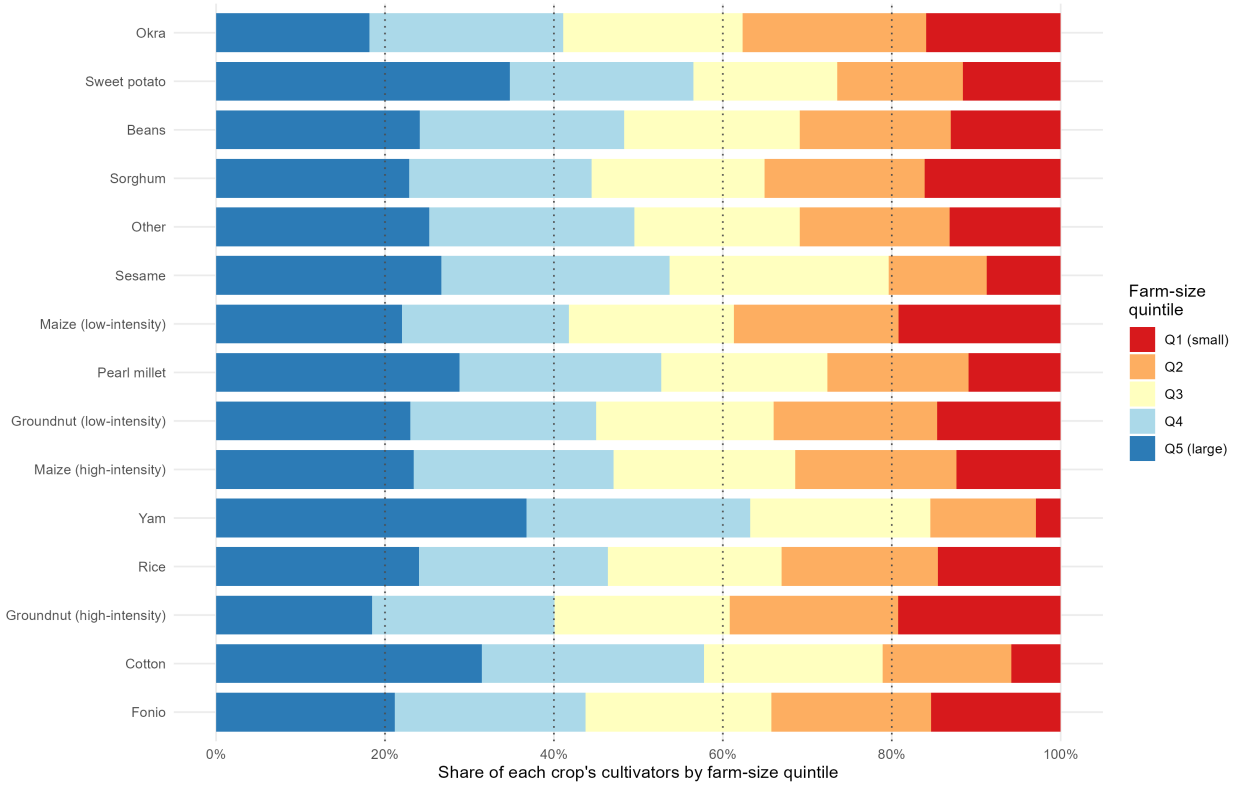


Notes: Cell (i, j) reports $|cultivators(i) \cap cultivators(j)|/|cultivators(i) \cup cultivators(j)|$ at baseline (2010). Crops ordered by estimated σ_c (top = lowest risk). The full conditional probability matrix $P(j \mid i)$ is included in the replication materials.

Figure J.3 addresses the same concern more directly by showing, for each crop, the share of its cultivators drawn from each baseline farm-size quintile. If a crop's cultivator pool were dominated by capable farmers assumed to be cultivating larger farms, its bar would be heavily weighted toward Q5 (largest farms). For every major crop except cotton, the share from the smallest-farm quintile (Q1) is between 11% and 19%, close to the uniform 20% that would obtain if cultivators were drawn at random from the population. Cotton is one crop with visible skew (Q1 = 5.8%, Q5 = 31.5%), reflecting the labor and processing requirements of cotton production; yam is another. The cotton-excluded robustness in Appendix G.1 confirms that this skew does not drive the headline portfolio standard deviation effect. Overall, farm-size selection does not

appear to be driving differential crop portfolio choices even though it is significant in Figure J.1 for every crop.

Figure J.3: Farm-size composition of each crop's cultivator pool

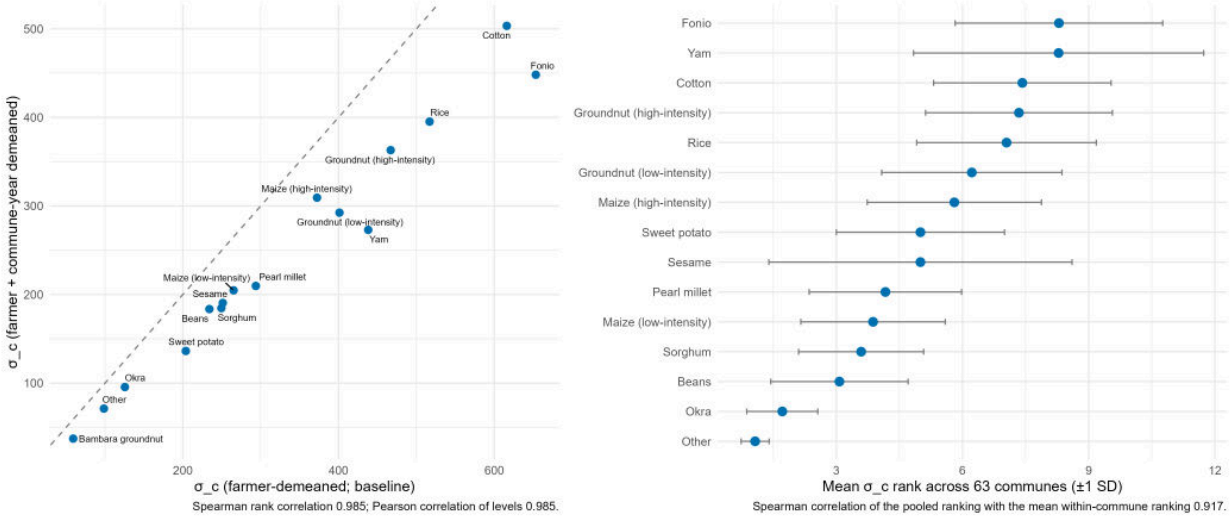


Notes: For each crop, the bar shows the share of its baseline cultivators drawn from each farm-size quintile, where quintiles are constructed from the full baseline distribution of total cultivated area. A uniform 20/20/20/20/20 split would be obtained if cultivators were drawn at random from the population. Crops ordered by estimated σ_c (top = lowest risk).

J.3: Geographic-shock contribution to σ_c ranking

A separate version of the selection concern operates at the geographic level rather than the household level: if risky agro-ecological zones cultivate only the safest crops, the cross-section of cultivators is geographically censored and the observed σ_c ranking could be driven by shared local shocks rather than by the underlying riskiness of each crop. Figure J.4 reports the same crop-level σ_c estimated under three demeaning specifications: (A) farmer-demeaned profit residuals (the baseline specification used throughout the paper), (B) farmer plus commune-by-year-demeaned residuals (which strip common geographic shocks), and (C) per-commune σ_c averaged across the 64 communes with at least 400 crop-year observations. The Spearman rank correlation between specifications A and B is 0.985; between A and the within-commune mean rank is 0.917.

Figure J.4: Crop σ_c ranking under alternative geographic demeaning



Notes: Left panel: scatter of σ_c under farmer-demeaning (x-axis) versus farmer plus commune-by-year-demeaning (y-axis). Diagonal reference line shows perfect agreement. Right panel: mean within-commune σ_c rank across 63 communes with ≥ 400 farmer-crop-year observations, ± 1 within-commune rank SD. Crop ordering on the y-axis is by mean within-commune rank.

The crop σ_c ranking is therefore not an artifact of geographic-shock heterogeneity. Common shocks are not driving the ordering: when they are stripped via commune-by-year demeaning, no crop's rank among the 16 moves by more than 2 positions. The within-commune analysis confirms that even within individual local-shock environments, the same crops appear at the top and bottom of the σ_c distribution.

Together, the three exhibits in this appendix indicate that selection on observable household characteristics, capability-based partitioning of the menu, and geographic-shock heterogeneity are each insufficient to drive the observed $(\vec{R}, \mathbf{C})$ rankings. The paper's causal claims are identified within the randomized experimental sample regardless: $(\vec{R}, \mathbf{C})$ characterize the menu of crop opportunities this sample's farmers actually face, and the welfare analysis is computed against that menu.

Appendix K: External Validation of Crop Moments against FAOSTAT Mali

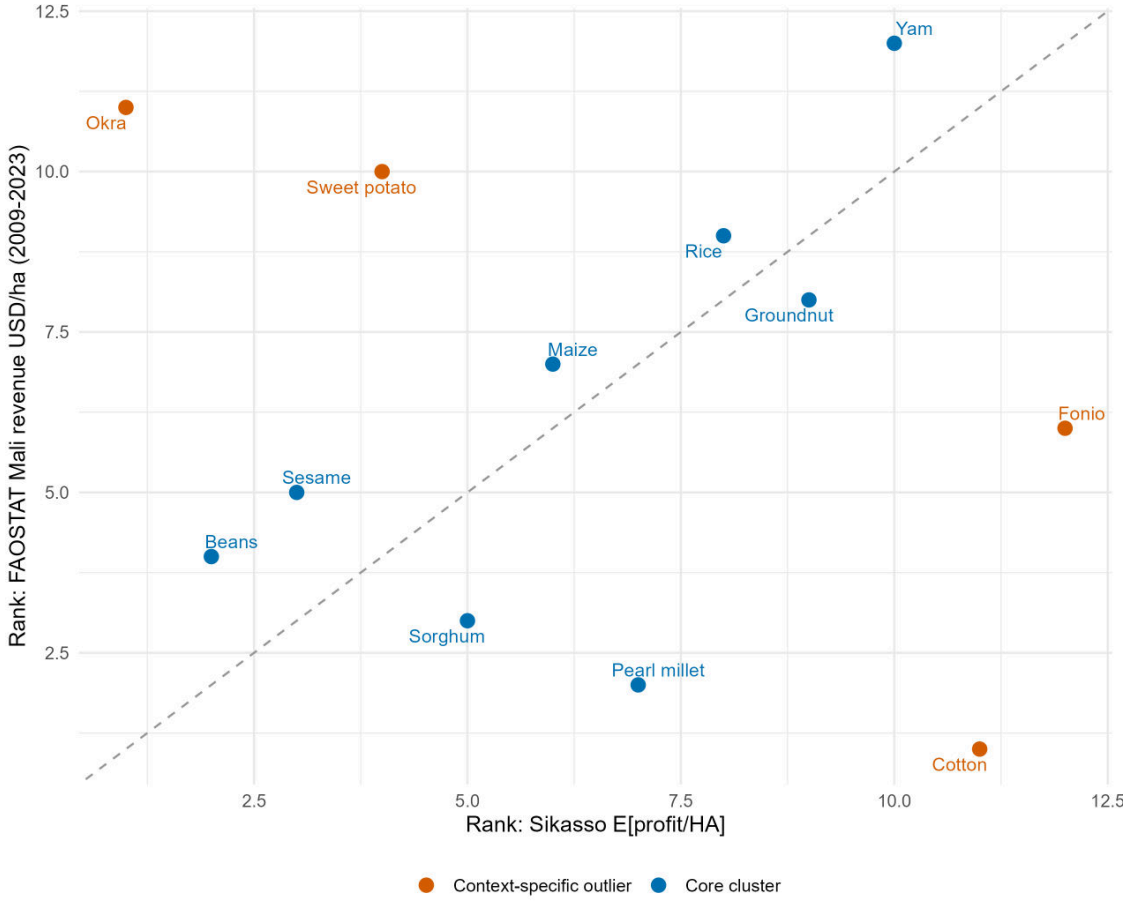
A natural concern with the estimated crop moments is that the four-year panel and within-farmer deviations yield noisy estimates that do not correspond to any external benchmark. I address this by comparing the Sikasso crop ranking on expected profit per hectare against an independent source: FAOSTAT national production statistics for Mali, converted to revenue per hectare using FAOSTAT's producer prices (USD/tonne).

For each crop in the sample, I compute annual national yield (kg/ha) and producer price (USD/tonne) from FAOSTAT for Mali over 2009–2023, chosen to overlap the RCT survey rounds.[26] I then construct a yield-price revenue measure $R_{kt} = Y_{kt} \times P_{kt}/1000$ in USD/ha for crop k in year t, and take the time-mean $\bar{R}_k$. Crops are ranked by $\bar{R}_k$ and compared to the Profit/HA ranking produced by the Sikasso sample estimates $\widehat{E}[\pi_k/\text{HA}]$ (Table B.6).

Figure K.1 plots the two rankings. Across all twelve crops the rank correlation is $\rho = -0.10$, and that near-zero figure is itself informative: the disagreement is not diffuse noise, but is concentrated in four specific crops. On the *core cluster* of 8 crops (yam, groundnut, rice, pearl millet, maize, sorghum, sesame, and beans) the Spearman correlation is $\rho = 0.67$, rising to $\rho = 0.86$ when pearl millet is additionally excluded. The remaining four crops (cotton, okra, sweet potato, and fonio) diverge for identifiable context-specific reasons summarized below.

[26] Results are similar for the longer 1994–2023 window; see `2D_FAOSTAT_Comparison_Long.csv` in the replication materials.

Figure K.1: Crop-Level Expected-Return Ranking: Sikasso Sample vs FAOSTAT Mali



Notes: Each point is a crop, ranked by its expected profit per hectare in the Sikasso sample (horizontal axis) and by FAOSTAT's national revenue per hectare for Mali, 2009–2023 (vertical axis); the dashed line marks equal ranks. The core cluster (blue, 8 crops) has a Spearman rank correlation of 0.67 (0.86 excluding pearl millet); across all crops it is -0.10. The four context-specific outliers (orange) diverge for the reasons given in the text.

Why the four outliers are expected to diverge

Cotton. FAOSTAT prices "seed cotton" at its unprocessed, pre-ginning value (roughly $143/ha of revenue in Mali in 2009–2023). In Sikasso, by contrast, growers sell through the CMDT contract system, which captures value further down the processing chain; the Sikasso sample reflects that integrated farm-gate profit and ranks cotton near the top in expected return.

Okra and sweet potato. FAOSTAT reports national yields of 12 and 16 tonnes/ha respectively, consistent with commercial horticultural operations, and the implied revenue exceeds $10,000/ha. In Sikasso, both crops are grown at subsistence-garden scale on small plots; the farmer-level profit estimates reflect that reality.

Fonio. Fonio appears at the top of the Sikasso ranking because the few farmers who grow it do so as a niche cash crop commanding a price premium. FAOSTAT's national average smooths this together with low-price Sahelian subsistence fonio, yielding a middling revenue rank.

The core-cluster rank correlation of 0.67–0.86 indicates that the estimated crop moments capture real cross-crop differences in profitability, not idiosyncratic noise from the four-year panel. The residual divergence is concentrated in four crops whose Sikasso-specific market structures differ systematically from national aggregates: a feature of the data, not a failure of estimation.

There are three caveats that should be discussed regarding this validation exercise. The first being that of comparing revenues and profits. The FAOSTAT only captures gross revenue while my sample estimates net out input costs. Cross-crop differences in input intensity create systematic wedges between the rankings. The second is that the FAOSTAT aggregates across Mali's two distinct agroecological zones (Sahelian north and Sudanian south) and Sikasso occupies only the latter. Finally, the FAOSTAT yield CVs reflect national-year shocks; the sample SD estimates reflect within-farmer year-to-year variation. The two statistical objects are not directly comparable, and I accordingly do not use FAOSTAT to validate the risk estimates in this exercise.