How Important are Investment Indivisibilities for Development? Experimental Evidence from Uganda*

Joseph P. Kaboski † Molly Lipscomb ‡ Virgiliu Midrigan § Carolyn Pelnik ¶

May 29, 2026

Abstract

Theoretically, indivisible investments can generate poverty traps and risk-loving behavior. We show experimentally that such traps shape households' medium-term micro-level investment and income dynamics. Still quantitatively they carry limited implications for long-run aggregate development. Given a choice between a safer, lower-payoff and a riskier, larger-payoff lottery, 27% chose the riskier option despite its lower expected value. Large lottery winners invested more in land and business durables, eventually increasing income; small winners' gains are transitory. Our model shows macro policy implications are muted: grant impacts dissipate, and financial deepening is weakened by inelastic land supply, though savings expansion remains somewhat effective.

*We thank Erin Byrne, Savita Diggs, Ambrose Donnelly, Shoshana Griffith, Suk Joon Kim, Laurene Lee, Jenny Macler, Regina Mannino, Grant Nelson, Virna Vidal Menezes, Priya Viswanathan, and Elizabeth Vranas for excellent research assistance. We thank participants for comments received on presentations of this research at numerous seminars and conferences. Special thanks to Cynthia Kinnan and Mark Rosenzweig for insightful discussions of the paper. We are grateful to the Kellogg Institute, the National Institute of Health, MasterCard Foundation, the Gates Foundation, and STEG for funding this project and research. This project was reviewed and approved by the Institutional Review Boards of the University of Notre Dame, the Joint Clinical Research Centre in Uganda, and Mildmay Uganda. The project was pre-registered in the AEA RCT Registry (AEARCTR-0002217).

†University of Notre Dame and NBER. Corresponding Author: Department of Economics, 3039 Nanovic Hall, Notre Dame, IN 46556, jkaboski@nd.edu

‡University of Virginia

§New York University and NBER

¶World Bank

1 Introduction

Poverty traps, whether household- or macro-level, are of primary importance to development economists and policy practitioners as a potential barrier to growth (e.g., Banerjee and Newman, 1993; Balboni, Bandiera, Burgess, Ghatak, and Heil, 2022). Poverty traps involve self-reinforcing dynamics, and a common ingredient is the presence of high-yield indivisible investment opportunities that increase income but can only be accessed by those who already have money. These vicious circles invite policy interventions such as one time grants, which enable poor households to escape immediately, or financial services, where households can either borrow (to make high-yield investments from which they repay loans with their increased income) or save out of their predicament. A key question is whether the presence of such indivisibility-driven poverty traps is important empirically, and whether policies can effectively address poverty traps and lead to long run development in their presence.

We combine a field experiment with a structural macro model to assess the empirical and quantitative importance of poverty traps for development in the context of semi-urban Uganda. Specifically, we use a field experiment to test the extent to which households will select into a riskier, lower expected value lottery for the chance to obtain a much larger payoff than a safer but lower payoff lottery in semi-urban Uganda where financial services are scant. The experiment enables us to empirically assess the rate at which participants take up the riskier lottery, the characteristics of those selecting the riskier lottery, the investment choices of lottery winners, and the effect of winning the lottery on income and consumption. We motivate the experiment with a model of indivisible investment opportunities in an environment without credit or interest-bearing savings. The model shows that households with a high-yield indivisible investment opportunity may be unable to escape poverty traps if they lack access to finance and sufficient wealth to fund the investment.

Although we find micro evidence consistent with indivisible investment-driven poverty traps, our macro model — together with key features of the data that discipline it — demonstrates that such poverty traps are not as important for long-run aggregate development as we might expect. Our macro argument rests on three pillars. First, we find considerable churning in the data with respect to assets and income, indicating that shocks are important and "traps" may not be as permanent as the term would imply. Although traps are persistent, eventually these shocks dominate so that neither initial conditions, nor the cash grants that alter them, are predicted to have impacts in the very long run. Second, we calibrate the model to match macro distributional moments drawn directly from the experiment — the distributions of income, assets, and lottery outcomes across groups — rather than micro regression coefficients. This disciplines aggregate predictions in a way that is consistent with modest long-run impacts even when individual treatment effects appear large. Third, the indivisibility is in land, which is inelastically supplied, so loans — another policy to address poverty traps — are largely ineffective, since they mostly drive up the price of land rather than expanding real investment. While the experiment uncovers micro poverty traps, the macro analysis shows that their importance is lessened in the long run and in general equilibrium.

Starting with theory, we show that although poor households are generally risk averse (e.g., Rosenzweig and Binswanger, 1993; Donovan, 2021), the presence of high-yield indivisible investment opportunities can induce some households to make risk-loving choices. Our model predicts that high productivity households with assets below what is needed for the indivisible investment are characterized by risk-loving behavior, due to high potential returns. They may also have higher initial savings, in an attempt to finance the indivisible investment good.

The lottery experiment then poses households with an actual choice between a lottery that pays out approximately $100 with 50% probability versus a lottery with lower expected value that pays out approximately $475 with 10% probability, allowing us to characterize who chooses the riskier lottery. The randomly selected winners receive the payout of their chosen lottery as a cash grant via mobile money. We follow up at 4 months, 18 months, and 6 years with a survey which allows us to empirically assess how investment goods vary by lottery choice. Given our semi-urban setting, where three-quarters of households operate a farm but more than half operate a (non-farm) business, neither the existence nor nature of indivisibile investment opportunies are *ex ante* obvious.

We find that a substantial share of the sample, 27 percent, does indeed prefer the riskier option. These riskier choices are more common among men, well-known to have lower underlying risk aversion (e.g., Meissner, Gassmann, Faure, and Schleich, 2023), but consistent with the model, the risky option is also preferred by larger, higher-saving, and investment-oriented households. Those choosing the riskier lottery have higher crop income, business assets, and total wealth and are more likely to report a desire for credit to make a business investment. We also find that only winners of the riskier, larger payoff lottery sustain higher levels of investment in indivisible goods at the second endline. Large lottery winners invest primarily in land, and land exhibits substantial capital gains. In contrast, small lottery winners make short-lived investments in small livestock (e.g., chickens) and business inventory (e.g., retail goods).

The model suggests that high-productivity households with new access to a large sum of money invest it and should be able to make large gains in income over time. In the short run (4 months after the lotteries pay out to winners), we find no significant impact on either consumption or income, by small or large lottery winners. In the medium run (18 months after payout), however, we find that large lottery winners have higher consumption, while winning the small grant still has no significant impact on income. In a long run follow-up survey (6 years after payout), we find a substantial increase in income for large lottery winners with some (imprecise) evidence of continued gains to consumption. This is consistent with findings from Balboni et al. (2022) in which a household's ability to leverage a randomized transfer of wealth to work its way out of poverty depends on initial levels of wealth. In sum, we find experimental evidence consistent with poverty traps being driven by the presence of indivisible investment opportunities, which are important for the medium-term micro dynamics of investment and income.

Next, we calibrate the model to reflect participants' actual lottery choices and the distributions of income, assets, and lottery outcomes across groups. We then use the calibrated model to simulate (i) long-run

predictions for the grants' impacts, and (ii) the impact of financial services expansion. The impacts of grants are persistent but not permanent, with half-lives of roughly 8–9 years, and impacts depend critically on the targeting of risk-seeking individuals. The impacts of relaxing financial frictions on aggregate development in the area depend critically on the elasticity of the supply of the investment good. If the lumpy investment is elastically supplied to a community (e.g., tractors for a farm or refrigerators for a business), then access to credit or savings expands the set of people able to invest, and the increase in demand expands the amount of real investment in the community. In this case, financial services can increase income, investment, and economic mobility. If instead the lumpy investment is inelastically supplied in the community (e.g., land), then increased demand from expanding financial services manifests largely in an increase in prices rather than real investment.

Finally, we use the experimental variation in demand to estimate that land is indeed inelastic. The model, together with the presence of substantial shocks, the fact that land is the dominant indivisible investment, and the fact that land is inelastically supplied even in semi-urban areas, leads us to conclude that neither cash grants nor credit can lead to long-run development in these areas.

The rest of the paper is organized as follows. Section 2 presents our simple motivating model. Section 3 summarizes our experimental methodology and data collection. Section 4 describes the types of households that select into the large lottery and the effect of lottery winnings on investment and the household budget in general. Finally, Section 5 calibrates the model to our data and simulates the macroeconomic consequences of credit injections and redistribution.

Contribution to the Literature

Our paper offers a "micro-to-macro" contribution to the literature on poverty traps, cash grants, and the impacts of financial services expansion. In particular, we add both new micro experimental evidence and a macro development perspective on the impacts of policies in the long run and general equilibrium. We also add experimental evidence that speaks to the the literature on broader aggregate importance of financial frictions for development.

First, we speak to research, both old and new, that presents empirical evidence of poverty traps and indivisibilities in developing economies. Balboni et al. (2022) examine the impact of a uniform livestock asset grant in Bangladesh and show that the impact is an s-shaped function of the initial assets of the recipient. On the micro side, our findings and methods complement theirs: we allow for an endogenous relationship between productivity and initial wealth, which would confound empirics based on the initial level of assets, so we focus instead on risk preference. Relatedly, Banerjee, Breza, Duflo, and Kinnan (2019) present empirical evidence that despite low average returns, microcredit does indeed break a poverty trap for "gung-ho" entrepreneurs, those with prior business experience who exhibit higher than average returns to business investment. In non-experimental settings, Lybbert et al. also examine livestock-based poverty traps empirically in a series of papers (e.g., Lybbert, Barrett, Desta, and Coppock, 2004; Carter and Lybbert,

2012), and they highlight the role of risk taking (Lybbert and Barrett, 2011). McKenzie and Woodruff (2006) present evidence for poverty traps with observational data on Mexican microenterprises. These papers emphasize business investment and livestock as indivisible investments driving poverty trap dynamics. To this literature, we contribute novel findings on the role of land in generating poverty traps.[1] Our primary contribution, however, is to show that even when poverty traps are present, the suggested policies may not be effective for long-run development.

One such policy is the giving of cash grants. Such grants have become a popular approach to identifying the marginal returns of capital for entrepreneurs (De Mel, McKenzie, and Woodruff, 2008, 2014) and the returns to poverty programs more generally (Blattman, Fiala, and Martinez, 2014; Blattman, Green, Jamison, Lehmann, and Annan, 2016; Haushofer and Shapiro, 2016, 2018; Egger, Haushofer, Miguel, Niehaus, and Walker, 2019). Recent work by Banerjee, Faye, Krueger, Niehaus, and Suri (2023) shows that cash grants of equivalent value have greater impact on revenues and enterprise ownership when delivered as a lump sum (rather than through multiple disbursements of smaller size). This is consistent with work by Casaburi and Macchiavello (2019), Herskowitz (2021), and Gertler, Higgins, Scott, and Seira (2023) who show demand for lumpiness, as evidenced by demand for less frequent but larger payments, the use of gambling to finance lumpy expenses, and greater take-up of prize-linked savings when there is the possibility of winning a large cash lottery, respectively. Other related work by Beaman, Karlan, Thuysbaert, and Udry (2023) shows evidence of selection into borrowing that reveals which farmers have high returns to capital. We combine experimental variation in grant payouts with a choice between a safer, lower-payoff lottery and a riskier, higher-payoff lottery to reveal which respondents are willing to make a risk loving choice in response to a potential poverty trap. Using the model, we demonstrate that such targeting to high productivity entrepreneurs facing liquidity constraints and indivisible investments is important but also that churning is an integral part of development, and this limits the long run impact of cash grants.[2] Financial expansion is another such policy. Our simulations of the aggregate impacts of financial frictions are an additional contribution. Methodologically, we contribute to an emerging literature using experiments in conjunction with macro development models (e.g., Buera et al., 2014; Bergquist, Faber, Fally, Hoelzlein, Miguel, and Rodriguez-Clare, 2019; Donovan, 2021; Lagakos, Mobarak, and Waugh, 2018; Buera, Kaboski, and Shin, 2020).[3] While these papers use the results of RCTs to discipline simulated impacts of scaled policies in the macroeconomy, our paper is novel in using an experiment to actually test an existing theory of macrode-

[1] Although land is in principle divisible, whether property rights are customary or formal, it is generally divided into discrete plots both in its use and any transactions and titling. In some cases, rental markets can be a way of overcoming investment indivisibilities (e.g., Bassi, Muoio, Porzio, Sen, and Tugume, 2022; Caunedo and Kala, 2021), but rental is less useful if capital gains are an important part of the return. Our findings on land also complement recent work by Acampora, Casaburi, and Willis (2022), who show that market imperfections lead to the under-utilization of land rentals and lower aggregate returns to land.

[2] Our evidence of churning offers empirical support for the assumptions in Buera, Kaboski, and Shin (2014) for its quantitative analysis of asset grants.

[3] See Buera et al. (2020) for a review of this literature.

velopment and quantify its implications.[4] We focus on financial services and illustrate that, even when indivisibilities can lead to micro-level poverty traps where financial services may seem particularly needed and powerful, their scaled impact depends critically on the elasticity of the supply of the relevant capital.

Finally, we add empirical and quantitative guidance to a macro literature on financial frictions and poverty traps. Banerjee and Newman (1993), Galor and Zeira (1993), and Piketty (1997) are examples of models with indivisibilities, where financial frictions lead to poverty traps. Later work showed that indivisible decisions, when embedded in quantitative models with intensive margins and mapped to the data, did not lead to macro multiplicities, but only micro poverty traps as in Buera (2009). This resulted in lower aggregate output (Buera, Kaboski, and Shin, 2011; Midrigan and Xu, 2014) and slower convergence (Buera and Shin, 2013). These theories all lead to demand for financing but also risk. We contribute to these findings by: (i) showing the empirical importance of both indivisibilities and risk-loving behavior linked together, (ii) using these empirical results to discipline their quantitative importance, and (iii) introducing an inelastically supplied capital good into the analysis.

2 Model and Motivation

In this section, we develop a model to demonstrate how the combination of high-yield indivisible investments and financial frictions can lead to both poverty traps and consequent risk-loving behavior. This provides motivation for our micro empirical experiment in the next section, and will be the basis of our macro simulations as well. The model features three key ingredients: multiple discrete capital sizes; a borrowing friction, and stochastic shocks to labor income, entrepreneurial productivity, and the discrete ability to operate capital. Although we focus on qualitative patterns of individual decisions here, the model is parametric, since it will be later used for quantitative analysis in Section 5.

2.1 Set Up

We discuss the preferences, productivity, and investment decisions of households in turn.

2.1.1 Preferences

Household i has time-additive, log preferences over consumption:

$$\hat{V}_{i0} = E_t \sum_{t=0}^{\infty} \beta^t \ln c_{it}.$$

We assume that a fraction $1 - p$ of them die each period and are replaced by an equal mass of newborn households with no initial wealth. Therefore, $\beta = \hat{\beta} p$ reflects time discounting, a product of the pure

[4]Moreover, another methodological innovation is our calibrating the model to match macro distributional moments drawn directly from the experiment — the distributions of income, assets, and lottery outcomes across groups — rather than micro regression coefficients. This connects the experimental variation directly to aggregate predictions in a way that disciplines the long-run and general equilibrium implications of the model.

discount factor $\hat{\beta}$ and the survival probability p. The expectation is over realizations of death, shocks to labor income, shocks to entrepreneurial ability, and the land ability shock introduced below.

2.1.2 Income

Labor income and entrepreneurial ability each have a permanent household-specific component and an idiosyncratic transitory shock. Specifically, we assume

$$\log e_{it} = \bar{e}_i + \varepsilon_{it}$$

where $\bar{e}_i$ is a permanent component and ε_{it} are idiosyncratic, independently- and identically-distributed (iid) innovations from $N\left(0, \sigma_e^2\right)$. We normalize $\bar{e} = 1$ without loss of generality. The permanent component $\bar{e}_i$ captures persistent differences in labor market productivity across households—arising, for example, from differences in education or occupational skill—while ε_{it} captures transitory fluctuations around this permanent level.

Entrepreneurial ability $\xi_{it} z_{it}$, is a product of two processes. The z_{it} term follows a process analogous to labor:

$$\log z_{it} = \bar{z}_i + v_{it}$$

where $\bar{z}_i$ is a permanent component and v_{it} are iid draws from $N\left(0, \sigma_z^2\right)$. Whereas v_{it} is transitory, ξ_{it} is a more persistent shock to productivity. The discrete shock, $\xi_{it} \in \{0, 1\}$, amounts to a productivity-driven capital-demand shock. Its evolution follows the Markov transition matrix

$$\begin{bmatrix} \pi_0 & 1 - \pi_0 \\ 1 - \pi_1 & \pi_1 \end{bmatrix},$$

where π_1 governs the persistence of being in the productive entrepreneurial state and π_0 governs the rate at which households that previously could not operate capital transition into doing so. A household with $\xi_{it} = 0$ earns no operating return from capital and therefore optimally holds none.

We therefore model three sources of heterogeneity in entrepreneurial productivity: permanent, $\bar{z}_i$, persistent, ξ_{it}, and transitory, v_{it}. The combined shock $\xi_{it} z_{it}$ introduces an additional source of risk: not only does productivity fluctuate each period, but households may also face spells during which capital generates no income at all. We assume households know their value of ξ_{it} at the end of period $t - 1$, before making their land holding decision for period t.

Capital income is a function of combined entrepreneurial ability, $z_{it}\xi_{it}$ and capital, k_{it}:

$$y_{it} = \xi_{it} z_{it} k_{it}^{\alpha} - \delta k_{it},$$

where δ denotes the maintenance cost of capital. Although in principle, capital is general, we will refer to it as "land", since this will be largely consistent with our empirical findings.

Indeed, despite the semi-urban location of our empirical application, we will find that agriculture is a primary use of such land, which warrants further comments: although we appear to abstract from two commonly-modeled features of agriculture, our model can largely capture them. First, although not explicitly *ex post*, realizations of our transient shocks, ν_{it}, lead to income shocks akin to shocks to yield (e.g., from weather).[5] In contrast to transient shocks, persistent shocks to ξ, will impact the demand for land, while $\bar{z}_i$ captures permanent agricultural ability that drives long run wealth distributions. We also abstract from aggregate shocks as a simplifying assumption, since aggregate uncertainty would complicate the solution algorithm. Such shocks are empirically relevant as noted in Section 4.4.3. But we note that aggregate shocks would only reinforce our key messages regarding long run development, as we explain in Section 5.4. Second, regarding labor, whether it is exogenously fixed to family size (as in LaFave and Thomas (2016); Foster and Rosenzweig (2022)) or endogenously hired at a fixed wage, ω, it is easily incorporated into our production function by interpreting the product $\xi_{it}z_{it}$ as labor-augmented productivity.[6]

2.1.3 Financial Environment and Investment Decisions

Households choose their liquid assets and land holdings. While liquid assets a_{it+1} earn an interest rate of r, land generates net operating income $\xi_{it}z_{it}k_{it}^{\alpha} - \delta k_{it}$. In principle, liquid wealth can be negative, but savings decisions are bounded below since borrowing is limited by the value of available land as collateral. We assume that households can borrow up to a fraction θ of this value:

$$a_{it+1} \geq -\theta P_t k_{it+1}.$$

A central feature of the model—and the key source of investment indivisibility—is that land can only take on a finite set of discrete values:

$$k_{it+1} \in \{0, k_1, k_2, \ldots, k_n\}.$$

The smallest non-zero size k_1 requires a minimum wealth threshold to finance, generating poverty trap dynamics similar to those in a model with a single indivisible investment. The existence of larger sizes $k_2, \ldots, k_n$ creates additional investment thresholds and a richer pattern of savings behavior around each. Because households with $\xi_{it} = 0$ hold no land, the investment decision is only active for households in the $\xi = 1$ state.

Denoting the price of land P, the household's budget constraint is:

$$c_{it} + a_{it+1} + P(k_{it+1} - k_{it}) = e_{it} + (1 + r)\, a_{it} + \xi_{it}z_{it}k_{it}^{\alpha} - \delta k_{it}. \tag{1}$$

[5] Modeling uncertainty as *ex ante* rather than *ex post* limits the availability of risky investments but is largely innocuous otherwise, since we abstract from default in our analysis of credit.

[6] To see this consider a more general production function, $B\xi_{it}\tilde{z}_{it}k_{it}^{\tilde{\alpha}}l_{it}^{\gamma}$. In the case of exogenous labor supply, we can define $z_{it} \equiv \tilde{z}_{it}l_{it}^{\gamma}$, $\alpha \equiv \tilde{\alpha}$, and $B \equiv 1$ to yield our production function. If instead labor is endogenously hired at spot rate, ω, solving for optimal labor and substituting in we need to define $z_{it} \equiv \tilde{z}_{it}^{1/(1-\gamma)}$, $\alpha \equiv \tilde{\alpha}/(1-\gamma)$, and $B \equiv (\gamma/\omega)^{\gamma/(\gamma-1)}/(1-\gamma)$ to recover our specification.

The household takes the price of land P as given.

2.2 Recursive Formulation

We reformulate the household's problem recursively. For convenience, we work with consumption-equivalent welfare (CEV), which is easier to interpret. The original utility $\hat{V}_{it}$ can be written as $\hat{V}_{it} = \frac{1}{1-\beta} \log V_{it}$, where V_{it} is the constant level of consumption that would make the household indifferent with their actual allocation. Under this transformation, the Bellman equation becomes:

$$V_{it} = \exp((1-\beta) \log c_{it} + \beta \, E_t \log V_{it+1}) \, .$$

The relevant endogenous state variable is total wealth,

$$w_{it} = (1+r) \, a_{it} + P k_{it} + e_{it} + \xi_{it} z_{it} k_{it}^{\alpha} - \delta k_{it},$$

which combines the return on liquid assets, the market value of land, labor income, and net operating income from land into a single cash-on-hand measure. The permanent components $\bar{e}_i$ and $\bar{z}_i$, drawn at birth, together with the entrepreneurial ability state ξ_{it}, complete the state of the household. The value function is therefore $V_{it} = V(w_{it}, \bar{e}_i, \bar{z}_i, \xi_{it})$, and the Bellman equation is:

$$V(w, \bar{e}_i, \bar{z}_i, \xi) = \max_{a', k'} \exp\big((1-\beta) \log c + \beta \, E \log V\big(w', \bar{e}_i, \bar{z}_i, \xi'\big)\big) \tag{2}$$

subject to the budget constraint

$$c + a' = w - P k', \tag{3}$$

the borrowing constraint

$$a' \geq -\theta P k', \tag{4}$$

and the law of motion

$$w' = (1+r) \, a' + P k' + \exp\big(\bar{e}_i + \varepsilon'\big) + \xi' \exp\big(\bar{z}_i + v'\big) \, \big(k'\big)^{\alpha} - \delta k'. \tag{5}$$

2.3 Model Results: Poverty Traps and Risk Loving Behavior

We first show how the model generates poverty traps and risk-loving lottery choices. We then summarize how the model motivates our experiment and empirics.

2.3.1 Optimal Choices and Poverty Traps

Poverty traps arise in the model due to indivisible land investment. Because households with $\xi_{it} = 0$ optimally hold no land, the investment decision is relevant only for households in the $\xi = 1$ state. The borrowing constraint $a' \geq -\theta P k'$ implies that the household can only choose to purchase a piece of land of

size k_j if their cash on hand w exceeds

$$w \geq \bar{w}_j = (1-\theta)Pk_j,$$

a condition necessary to guarantee positive consumption. Since investments in land are risky, however, households generally require a greater amount of wealth before purchasing a given plot size, in order to avoid fluctuations in consumption.

We illustrate the household's optimal decision rules in Figure 2. The four panels show the CEV value function (upper-left), land choice (upper-right), consumption (lower-left), and savings (lower-right), each plotted against total wealth w. Two lines are shown in every panel: one for $\xi = 0$ and another for $\xi = 1$. (All parameters follow our later empirical calibration in Section 5.)

The upper-left panel shows that the CEV value function is increasing in w for both ξ states, but the $\xi = 1$ household always attains a higher value, reflecting the additional option to invest in land. Notice that the value of an $\xi = 1$ agent has kinks at the investment thresholds where the land choice jumps, while the value of an $\xi = 0$ agent is smooth and concave. Note the convex segments of the $\xi = 1$ value function just below each threshold.

The upper-right panel shows that the $\xi = 0$ household never acquires land. The $\xi = 1$ household's land choice is a staircase: zero below the first threshold, then a discrete jump to k_1 upon crossing it, and further jumps to k_2 and k_3 at the subsequent thresholds. Households at every level of w would like to hold the (risk-adjusted) highest-return land size attainable, but poorer households are prevented from doing so by the borrowing constraint.

Looking at the lower panels, consumption (left) and savings (right) display strikingly different patterns across the two ξ states. For the $\xi = 0$ household, both policies are smooth and roughly monotone in w: without the option to invest in land, the household simply smooths consumption and saves at the market interest rate. For the $\xi = 1$ household, both policies are highly non-monotonic, with sharp dips at each investment threshold. At each threshold, the household commits nearly all available resources to land, accepting low values of consumption and running down liquid savings–because the return from investment is so high. As w rises beyond a threshold, consumption recovers and grows with wealth until the next threshold is approached, at which point the cycle repeats. This non-monotonicity in consumption with indivisible investment was emphasized by Kaboski and Townsend (2011), and such households in advanced economies were later coined the "wealthy hand-to-mouth" by Kaplan and Violante (2014) and Kaplan, Violante, and Weidner (2014).

Reflecting all these features, the gap in the value function between $\xi = 1$ and $\xi = 0$ is smallest at the lowest wealth levels, where neither household can invest, and larger at higher wealth levels, where the productive household can acquire larger and larger land holdings. The convex segments of the $\xi = 1$ value function—one for each threshold—are the source of the risk-loving behavior we turn to next.

2.3.2 Lotteries and Risk

The convex regions of the $\xi = 1$ value function lead to a preference for risk that motivates our empirical study. In the lottery experiments of the next section, we offer a choice between a risky (R) and a safe (S) lottery.[7] The safe lottery pays $\Delta^S = 1.29\bar{Y}$ with probability $p^S = 0.5$, where $\bar{Y}$ is median income, while the risky lottery pays $\Delta^R = 6.27\bar{Y}$ with probability $p^R = 0.1$. Note that although neither lottery can yield losses, the R lottery is much riskier in the sense of Rothschild and Stiglitz (1976), with a much higher mean-adjusted variance (16.028 vs. 0.393). It also has a slightly lower mean payoff (0.627 vs. 0.645).

In the CEV formulation, the values of the immediate lotteries are:

$$V^j(w, \bar{e}_i, \bar{z}_i, \xi) = \exp\Big(p^j \log V\Big(w + \Delta^j, \bar{e}_i, \bar{z}_i, \xi\Big) + (1 - p^j) \log V(w, \bar{e}_i, \bar{z}_i, \xi)\Big)\,.$$

for $j = R, S$. As in our empirical experiment, the risky lottery has a lower expected payout ($p^R\Delta^R < p^S\Delta^S$), so standard risk aversion in preferences pushes households toward the safer lottery. However, the indivisible investment opportunity and the borrowing constraint make the lottery preference depend on total wealth.

Figure 2 shows the difference $V^R - V^S$ (upper panels) and the optimal land choice k' (lower panels), each plotted against w, for a household with low $\bar{z}$ (left column) and high $\bar{z}$ (right column). For the low-$\bar{z}$ household (left), the difference is negative at all wealth levels: though the excess return to land is positive and the household would invest if it had enough resources, the return is too low to make gambling worthwhile, and the household always prefers the safe lottery. For the high-$\bar{z}$ household (right), the pattern is strikingly different: the household prefers the risky lottery in certain wealth ranges, shaded in gray. These bands align precisely with the investment thresholds visible in the lower-right panel: the household prefers the risky lottery in wealth ranges just below each step in the land choice staircase.

The economic logic is the same as in the single-threshold case, now repeated across each step of the investment ladder. When the household's wealth is just below a threshold, the safe payout (Δ^S) is insufficient to enable investment while the risky payout (Δ^R) is large enough to close the gap, making the gamble attractive despite its lower probability and expected value. Once wealth rises above a threshold, the land is already acquired and the value function is locally concave, so the safe lottery is again preferred. At wealth levels far below the next threshold, the investment opportunity is too distant to influence lottery choice, and the locally concave value function again makes the safe lottery optimal.

The result that risk-loving behavior is concentrated near investment thresholds and absent for low-$\bar{z}$ households ties the lottery choice directly to high-return investment motives rather than to any intrinsic taste for risk. With multiple land sizes, these threshold effects repeat at each rung of the investment ladder, generating a richer cross-sectional pattern of risk preferences that varies with both wealth and permanent entrepreneurial ability.

[7] We abstract from a separate dimension of our experiment described in Section 3.1 that involved a brief delay in proceeds because it did not yield results of great interest. Our model and simulations nevertheless account for the choices of this one-period delay.

2.3.3 Predictions of the Model

In summary, the model shows a connection between high-yield, indivisible investments, financial frictions, and poverty traps. The model motivates our experimental design in the next section and also leads to testable predictions that guide our empirics:

- high productivity households may have higher savings (shown in Figure 1, lower-right panel) to eventually finance indivisible investment opportunities (Figure 1, upper right panel) ;
- risk-loving lotteries can be chosen to help finance indivisible investments (Figure 2, upper-right panel);
- those choosing risk-loving lotteries have high entrepreneurial productivity (Figure 2, upper-right vs. upper-left panels); and
- these investments lead to persistent increases in income (right hand side of Equation 1) and ultimately consumption (Figure 1, lower-left panel).

3 Field Experiment

Our experiment presents participants with actual lotteries that mirror the theoretical lotteries of the previous section. This section provides details of the design and implementation of our lottery experiment as well as our data collection.

3.1 Experimental Design

To draw our sample, we worked with a prominent microfinance bank that was hoping to expand services. They identified three geographically dispersed, underbanked districts where marginal populations lacked financial services and they were considering expanding their services: Ntungamo, Ibanda, and Kagadi.[8] All three are district capitals with populations of roughly 20,000-30,000. The surveyed neighborhoods are best described as semi-urban and rural. Using a neighborhood census of each of the targeted neighborhoods, we randomly selected a sample of 1,048 participants, each from a distinct household.[9] We oversampled entrepreneurs and those who lacked formal financial services (i.e., had no formal loan or savings account at the time). All three districts are evenly represented in the sample: 350 participants come from Ntungamo, 349 from Ibanda, and 349 from Kagadi.

The timing of the experiment was as follows. In October and November 2015, we conducted a baseline interview of households including questions about their household demographics, income, consumption, agricultural and business activities, assets, borrowing and savings, and hypothetical preferences over income, investment, risk and willingness to delay payment for higher returns. At the end of the survey, all

[8] We piloted a related project in Mpigi in order to evaluate our survey instruments and other protocols.

[9] Neighborhoods were randomly selected by placing a grid across each city and randomly selecting gridpoints overlaid on a map of each city. Gridpoints that fell in manufacturing zones, parks, or other unpopulated areas were omitted. We then conducted a census of the neighborhood around gridpoints. From the census, we selected the sample, stratifying by use of formal financial services, expressed desire for a savings account, and occupation (entrepreneur, salaried employee, or farmer). We also stratified over gender and whether or not the recipient was a head of household.

households were offered a free, zero-interest formal savings account at a microfinance bank with a local branch. In February and March of 2017, approximately 16 months after participants received their savings accounts, we conducted the midline survey, where we resurveyed participants about their updated responses to select baseline questions and then implemented the cash lottery experiment described below. We then conducted a comprehensive first endline survey in June and July of 2017. We surveyed those who had chosen delayed payment one month after surveying the non-delayed group, so that all participants were surveyed approximately 4 months after receiving their cash grants. The first endline survey consisted of similar questions to the baseline, but included additional questions about how recipients had used their grant money. We conducted a second endline survey, meant to track medium-run household outcomes and observe changes in grant effects over time, in September and November 2018, approximately 18 months after cash grant receipt. Finally, we conducted a third endline survey to track long-run outcomes in June and July 2023, 6 years after grant receipt.

The experiment at midline asked the participant to choose between the following two lotteries: (i) a 50 percent chance of winning a grant of 350,000 Ugandan shillings or UGX ($97 based on an exchange rate of 3,600 UGX = 1 USD in 2017) or (ii) a 10 percent chance of a grant of 1,700,000 UGX ($472).[10] The sizes of grants were chosen based on baseline questions about desired investment amounts and demand for credit.[11] As in the model of the previous section, the second lottery has much higher risk in a Rothschild-Stiglitz sense, and a slightly lower expected value ($47.20 vs. $48.50). We visually primed participants through "practice" lotteries in order to assist them in understanding the probabilities and the lottery. Finally, we gave participants the choice to receive their grants (conditional on winning) via mobile money either the following day or to delay payment for 30 days in return for 3 percent interest. Our use of mobile money and a minimum of one-day delay was designed to limit differential perceived risk in whether they would actually receive funds since no one received cash immediately in hand.[12] After asking questions about how they would use their grants, we ran the lotteries using random number generators on tablets. Participants learned their outcomes immediately.

The choice of the small, less risky lottery with no delay was the most common, with 765 people (73% of the full sample) choosing the small grant. However, despite its riskiness and lower expected payout, 283 people (27% of the full sample) chose the lottery for the larger grant. This is quite a powerful result since the

[10] The effect of winning the lottery, thus, is identified conditional on lottery choice — whether the respondent selected into the small or large lottery. We do not cross-randomize participants into winning the lottery not of their choice due to budget limitations. Any additional large grants to small lottery winners would have required lowering the number of large grant chooser who we randomized to win. We address cross-randomizing grants via the quantitative model in Section 5.3.

[11] Specifically, we asked at baseline whether they would invest if they had access to credit, and how much they would need to make their desired investment. The sizes of the grants approximately match the 10^{th} and 75^{th} percentile responses, the large grant is 4.86 times larger than the small grant

[12] Recall that indivisibilities impact savings behavior (Figure 1) at wealth levels of agents choosing the small or large lottery (Figure 2). Of those who chose the small lottery, 144 people (14% of the full sample) chose to delay the payment by a month in return for a larger payment, while the remaining 621 (59%) chose to receive the payment the next day. Of those who chose the large lottery, 78 participants (28% of the high risk sample, 7% of the full sample) were patient and willing to wait a month for a larger payout, while 205 (20% of the full sample) were impatient. The high level of impatience is perhaps surprising given the high foregone rate of return (43%, annualized and compounded) and non-negligible absolute return (equivalent to 2-6 days of income for the median household, depending on the size of the grant) and will require high discount rates in order to reconcile.

large lottery had a lower expected value and risk aversion tends to increase in lottery size (Holt and Laury, 2002). The small lottery payout was equivalent to roughly 1 month's average household income, while the large lottery equated to 4 – 5 months of average household income. As shown in the theory of high-yield indivisible investments in Section 2, participants may be willing to accept higher risk lotteries in order to enable a large, high return investment.

The sizable fraction of people willing to choose a lower expected value, large lottery is consistent with the model's prediction driven by the presence of a high yield lumpy investment (and even quantitatively the 27% corresponds closely to our choosing the large lottery amount to be the 75th percentile of stated desired investment in the baseline. See Footnote 11). We call these people and their choice "risk loving" throughout, though it is important to note our interpretation through the lens of the model: risk-averse agents select the risky lottery because of potential large investment payoffs not because of a relative preference for risk.

Figure 3 presents the number of lottery winners in each category. We present the number of winners as a percentage of the people who made the choices. Of the 765 who chose the small lottery, 373 won the lottery for a win rate of 49%. Of the 283 who chose the large lottery, 85 won, for a win rate of 30%. Three points merit attention. First, we knew that the larger lottery (i.e., "risk loving" choice) was less likely to pay out and correctly anticipated that it would also be less popular. So in order to increase the sample size of winners, we increased the actual probability of winning the larger grant to 30% rather than the 10% probability communicated to participants. This was not disclosed to either the participants or enumerators performing the experiment. Second, one could be concerned about the possibility that the higher payout rate was inferred by participants, and in Appendix A.1 we conduct several tests that reassuringly suggest this was not the case. There is no statistically significant difference in the probability that participants choose the large lottery in the later days of surveying within each district (relative to earlier days of surveying), which we might expect if participants watched neighbors and friends win the large lottery and inferred its higher-than-stated expected value. Rather, our findings in Appendix A.1 are consistent with participants maintaining the same evaluation of the large lottery's expected value over time. Third, because of budget and sample size limitations, we did not randomize any participants into winning the small grant if they lost the large grant, or into winning the large grant if they lost the small grant. We instead analyze this through the model in Section 5.2.

The lotteries themselves provide our experimental variation. The lottery *choices* are of course endogenous, so we control for lottery choice in all specifications. Conditional on the lottery choice, the lottery *outcomes* are random. As discussed below in Section 3.2, we find that pre-intervention levels of key outcomes are well balanced between those winning versus losing the small lottery, as well as those winning versus losing the large lottery.

3.2 Data

We collected data over five survey waves: baseline, midline, first endline, second endline, and third endline. Each survey wave includes modules on income, consumption, agricultural activities, non-farm business activities, assets, borrowing and savings.[13] For several of the financial outcomes, we collect data in a multi-step process in order to improve measurement.[14]

Using these data, we look first at the use of the grant on investments: consumptive assets (home durables), productive divisible investments (small livestock, agricultural tools and business inventory), and productive indivisible investments (large livestock and ploughs, durable business assets, and land). Next we observe the impact of the grant on the other components of a household budget: consumption, savings, income, and net credit. The distinction between flows and stocks is important from a theoretical standpoint, as well as for interpretation through the lens of household's budget constraint and for aggregating categories. We measure income as a flow, and we measure savings, assets (business, agricultural, and land), and net credit as stocks at the time of survey. Income is the monthly flow of realized household income (crop profit, livestock profit, business profit, wages, and remittances). Savings and assets reflect current total levels at each respective endline. Net credit is current debt less any lending to others that the household expects to be paid back. At the first and second endline, we measure consumption as total spending on household expenditures (food, transportation, fuel, airtime, and any irregular expenses, like school, hospital or marriage fees) since the midline plus value of home durables. We combine household expenditures, which are collected as flows, and home durables, which are collected as stocks, together into this single measure of total spending between the midline and the respective endline. This differentiates household durable spending from spending on *productive* assets, which we include in the investment category because they may lead to positive income returns. Data on durables were not collected at midline, so we simply add the stock of durables at endline to the sum of consumption flows realized between midline and endline. This should lead to no issues with our application of a household budget constraint, as described in Section 4, as long as home durables were balanced at midline, as they were at baseline per Tables 1 and 2. At the third endline, we analyze impacts on the flow of weekly consumption separately from the stock of home durables, given that we are more concerned with understanding how potential long-run impacts might manifest, rather than detecting the impact of the grant on household budget components.

We also measure land values as estimated by the participant in all three endline waves. In the third

[13] The midline survey, however, is somewhat briefer than other surveys and does not collect updated information on all of the household assets included in the baseline survey. Specifically, we do not update the value of home durables, land, livestock and other agricultural assets from baseline to midline. Tables 1 and 2 depict balance on the most recent pre-intervention measurement of each outcome.

[14] First, we use questions about subcategories in order to prime the respondent about the many different specific sources of income, expenses, assets, and liabilities that may be relevant. Second, we summed the subcategories to get the total value and confirmed with the respondent whether the aggregate of the category reflects their perception. For example, income is collected as "typical monthly income" (revenues net of costs) using detailed questions about typical monthly revenues on subcategories — i.e., business, crop, livestock, and labor income — and typical costs, normalized to a monthly frequency, as well as direct questions on aggregate income and following up to see whether the participant viewed the aggregate or sum of the components as a better measure. We collect the respondent's income separately from household income because the former may be measured with less noise. The Data Appendix (Appendix A) describes these steps and measures in more detail.

endline, we add retrospective questions on purchases, sales, and investment (collected separately from total value). As with many semi-urban areas in developing countries, land values in Uganda increased significantly over time (e.g., Wineman and Jayne (2018); Gochberg (2021)). Because there was significant overall appreciation in land, we estimate rates of appreciation by imputing capital gains using *control* households. That is, we calculate the average increase in land values over time using the land values of control households.[15] On average, land values appreciate an estimated 41% over the 18 months between midline and second endline, or 2% per month (see Table A.2 and Appendix A for more detail on capital gains and their calculation).

Land value appreciation may either be a result of general local price increases, average general investment in land, or spillovers from the experiment itself, i.e., increased demand for land as a result of increased availability of capital due to the lottery. To estimate the extent to which the experiment itself may have impacted land prices, we calculate the percent of the increase in land value that can be attributed to nearby households receiving a grant and purchasing land in Appendix B. We use a two-stage least squares approach that instruments the change in an area's land values with the number of grants disbursed in that area. Table B.1 shows the first stage results, and Table B.2 shows the second stage results. As detailed in Appendix B, we find that land appreciated locally (within 1 mile) roughly 20% over the entire 18 month period between the midline and second endline as a result of the grants themselves, or approximately (20%/41%=)49% of the total land appreciation during the period. This price increase will play a pivotal role in interpreting our aggregate simulations, as it implies that the supply of improved land is relatively inelastic.

Tables 1 and 2 describe pre-intervention sample characteristics and confirm balance between the treatment (grant recipients) and control (non-recipients) for both the small and large lotteries, respectively. The average monthly income at midline is 362,000 UGX, or about $100 USD, and households average five members. Households are therefore quite poor, and even the small grants are sizable relative to income. Relevant to the demand for indivisible investments, more than 50% operate a non-farm business, while over 70% have a farm.[16] Financial services are not widely used – savings and debt levels are small. For example, the savings of large lottery winners are less than 2% of the value of their total assets, and the numbers are smaller for debt and for the other groups.

Comparing treatment and control, balance is quite good. For the 26 outcomes in the two tables, we would expect roughly 1 variable per table to be significant at a 5 percent level based on type I errors, but indeed none are statistically significant at the 5 percent level, and only 1 - 2 per table are significant at the 10 percent level. In particular, households winning the small lottery have slightly fewer adult females (1.1 vs 1.2, $p < 0.10$) and have higher non-stock business assets (421,000 UGX vs 282,000 UGX, $p < 0.10$) than the small lottery control. Large lottery winners are more likely than the large lottery control to have education

[15] Note that, in principle the overall increase in land values might not be entirely attributable to capital gains. It may be that, on average, all households in these semi-urban areas of Uganda are investing in their land (thereby increasing the land value). This would over-state the total level of capital gains over the period, but it would not affect our estimates of gains in land value for treatment households *relative* to control households.

[16] These are not mutually exclusive – most households have more than one source of income – and the reliance on both non-agricultural and agricultural income reflects the semi-urban environment.

beyond primary school (0.31 vs 0.20, $p < 0.10$). This is part of our motivation for controlling for baseline and midline levels in our endline analyses, as well as other economic and demographic characteristics such as midline income and the number of adult males, adult females, and children. We re-survey 74% of the original 1,048 respondents at all three endlines and show that our key findings are robust to various assumptions about the attrited sample in Appendix C.

4 Empirical Results

The model predicts that high productivity households who have large, indivisible investment opportunities are more likely to choose a lower expected value but larger potential payout lottery and that these investments should increase their income over time. We therefore present our empirical results (and methods) in three steps, starting with (endogenous) lottery choices, how winnings affect investment, and finally how they impact the broader household budget through consumption, income, and savings. Although we focus on the empirics through the lens of the model, the richness of the data also allows us to delve into deeper detail than the model's predictions.

4.1 Who Chooses the Large Lottery?

We examine significant midline predictors of lottery choice, and, consistent with the model, we find that participants choosing the large lottery tend to be those with higher income, with a self-reported desire to invest, with higher wealth, and whose wealth has increased quickly since the baseline. Within these results, we also find a special role for agriculture, including higher agricultural income and larger family size, an indicator of effective family productivity in agriculture.

We start with the following regression:

$$D_i^m = \beta P_i^m \left(+ \gamma X_i^m \right) + \varepsilon_i, \tag{6}$$

where D_i^m is the decision of household i at midline m (e.g., a dummy for the risky choice, a dummy for the impatient choice) and P_i^m are the midline predictors on which we test for significant differences between those choosing the small lottery relative to those choosing the large lottery. We run these tests with and without demographic controls X_i^m, a vector of household-level characteristics including household income at midline, female, household head status, age, age^2, number of adult females, number of adult males, number of children, district fixed effects, and patience (when testing for differences in the risk decision). We also use lasso (Tibshirani, 1996), to select the most important predictors of the household's lottery choice from among 160 covariates given a parameter, chosen using adaptive lasso, that penalizes additional model complexity.

Table 3 presents statistically significant predictors of the large lottery choice in the bivariate OLS regression that estimates the association between each predictor of interest and lottery choice without demo-

graphic controls X_i^m. Predictors also selected by the much higher dimension lasso regression are designated with an asterisk.[17] The first result of note is that risk choice does not appear to be random, which is comforting in that households were not purely choosing randomly (as might be the case, if they didn't understand the tradeoffs). In particular, we find that the lottery choice is correlated with proxies for underlying risk preference but also economic variables broadly consistent with the theory.

For example – and focusing first on demographics – we see that those choosing the large lottery tend to be men, who are well-known to have a greater preference for risk (Croson and Gneezy, 2009). They also come from larger households with more children and adult males, consistent with high effective productivity in agriculture, since family size is linked to farm labor supply (LaFave and Thomas, 2016; Foster and Rosenzweig, 2022).

More generally, we find the model's prediction that high productivity, investment-oriented (i.e., high ξz) households choose the riskier lottery bears out in the data. Lacking direct productivity measures, we use income and consumption as proxies. Those who select the large lottery have significantly higher incomes, especially crop income (both in absolute terms and as a fraction of their total income), significantly higher (13 percent) consumption, and significantly larger (14 additional log points) growth in consumption from baseline to midline.[18] Crop income, in absolute level and per adult equivalent, and growth in consumption, business income, and total monthly income are also selected by lasso.[19] Consistent with the model's prediction that investment-oriented households will prefer the large lottery, a significantly greater fraction (6 percentage points) of the participants who choose the large lottery report that they could increase their income if they had access to credit. Moreover, a significantly larger fraction (4 percentage points) want to invest an amount greater than $100, the size of the small grant, and a significantly larger fraction (7 percentage points) indicate that they would use credit for business investment. These predictors are significant at the 5% level in the OLS regression and all selected by lasso. Although many households might report a desire for credit generally, the significant correlation of these investment-oriented variables with the large lottery choice is consistent with them proxying for households in the high-productivity entrepreneurial state — those for whom the large grant specifically would enable a threshold-crossing investment that generates a positive return, rather than households currently in the non-productive state for whom investment would yield no return regardless of grant size. Thus, consistent with the model, it appears that the demand for the large lottery could be driven by a desire to invest, at least for some.

[17] The full list of covariates on which we test for significant differences between those choosing the small lottery versus the large lottery is in Table D.1 in the Appendix, and those selected by lasso are in Table D.3. The point estimates in both Table 3 and Table D.1 are from the unconditional specification of Equation 6, i.e. without X_i^m. The estimates conditional on X_i^m are available by request.

[18] On average, incomes increased significantly between baseline and midline faster than one would anticipate from pure trend growth. This may be due to aggregate fluctuations, seasonal fluctuations in measured monthly income or seasonal variations in the components of income (crop income was especially high, while business income was lower).

[19] A natural question is whether the financial predictors are ultimately driven by the demographics. For example, if men are both more likely to choose the large lottery and have a higher propensity to invest or accumulate assets, then perhaps the accumulation of assets is no longer predictive after controlling for gender. We analyze this by looking at the same predictions but controlling for household demographics. The findings are largely robust. One exception is the level of savings is no longer a predictor of selecting the large lottery once income is added as a control, though growth in savings from baseline to midline, as well as the level of business assets and wealth remain statistically significantly correlated with lottery choice. Table available on request.

Lastly, those selecting the large lottery experience *faster* increases in their savings from baseline to midline and are wealthier, especially in terms of business assets and land. Between baseline and midline, those selecting the large lottery experience significantly larger increases in total wealth (70 additional log points) and total savings (86 additional log points).[20] Thus, the accumulation of assets, also selected by lasso, seems to be a strong predictor of the choice of the large lottery.

We conclude that risk choices are not random, and reflect economic realities broadly consistent with the model, and the land results are of special interest.[21] The full predictive power of these explanatory variables is certainly not complete, however, and so we don't preclude that some randomness is present in the decisions.

It is worthwhile comparing these actual risk choices at midline with two *hypothetical* risk decisions involving investment collected at baseline. In contrast to the actual lottery choice at midline, the riskier options in these hypothetical survey questions involve the risk of absolute losses but with less variance and positive expected gains for the riskier option.[22] That is, neither decision involves actual risk-loving behavior, as the midline choice does. We find higher willingness to take these risks, and the percentage is higher for the expected higher payoff. As with the actual lottery choice, regressing these risk choice measures on demographics and economic states yields similar relationships though with slightly different variables showing up in the lasso specification (in Table D.3 in the Appendix). Namely, demographics predict risk choice in sensible ways, and those with higher income, wealth, and consumption are more willing to take risks. In all three risk choices, we therefore interpret that willingness to take risk depends on both underlying preferences (proxied by demographics) and by economic conditions. The correlation between baseline and midline choices is low, however, consistent with stochastically changing economic situations, as in the model.

4.2 Measuring the Impact of Grants

We use OLS regressions to estimate the impact of the large and small grants, conditional on the participant's choice between the small and large lottery. For each household i in district d, we regress:

$$Y_{id}^{e} = \beta_0 + \beta_1 Win_i^m + \beta_2 Win_i^m * D_i^m + \beta_3 D_i^m + \rho Y_i^b + \gamma X_i^m + \lambda_d + \varepsilon_{id} \tag{7}$$

where b denotes baseline, m denotes midline, and e denotes endline.[23] D_i^m reflects the participant's lottery selection, where $D_i^m = 1$ if the they selected the large lottery, $Win_i^m = 1$ if they won either lottery (small or

[20] These changes are large because of the very low base levels of wealth in the sample: 25% of the sample has 0 wealth at baseline. Before taking logs, we set any negative or 0 values equal to the minimum value + 1 Ugandan shilling.

[21] Though the theory is more ambiguous, we also test for significant predictors of the impatient versus patient, but find few differential characteristics. We control for patience in the remainder of our analyses, but do not interact the lottery outcome with time preference.

[22] Specifically, the first question asks whether the participant would make a risky investment option that yields either 10,000 UGX or 300,000 UGX with equal probability (expected value of 155,000 UGX) at a cost of 100,000 UGX. A follow up question asks the same payoffs but at a cost of 150,000 UGX.

[23] Recall that both baseline and midline take place pre-intervention, as the midline survey concludes with the household's lottery choice. We use e to generally denote the endline survey waves, but results are reported separately for each endline.

large), and their interaction, $Win_i^m * D_i^m$, is the additional effect of winning the large lottery. Thus, in this specification, β_1 identifies the (total) effect of winning the small grant, β_2 identifies the additional effect of winning the large grant relative to the small grant, and $\beta_1 + \beta_2$ identifies the total effect of winning the large grant. We conduct F tests on the sum $\beta_1 + \beta_2$. Critically, β_1 and β_2 are estimated conditional on lottery selection, as captured in D_i^m. (In our regression tables, *won lottery* corresponds to Win_i^m, *won large lottery* corresponds to the interaction term $Win_i^m * D_i^m$, where large lottery winners are those who both won their lottery and selected the large lottery, and the coefficient on D_i^m is not shown due to space constraints.)

In this regression model, λ_d are district fixed effects and X_i^m is a vector of household-level demographic controls: patience, household income at midline, age, age^2, gender, household head status, number of adult females, number of adult males, and number of children. We examine impacts on endline levels conditional on baseline levels (as denoted by Y_i^b), consistent with the prescriptions in Bruhn and McKenzie (2009).[24] We winsorize to the 5th and 95th percentiles to ensure that results are not driven by outliers — a concern given the high degree of skewness in the income and asset distributions — and inflate/deflate all variables to constant UGX at the time of the midline (2017). We apply multiple hypothesis corrections in Appendix E. In the following subsections, we examine the impacts of winnings first on stock variables (assets and durables), to understand how recipients spent grants, and subsequently on other components of the household budget — both flow variables (consumption, income) and stock variables (savings and credit).

4.3 Impacts on Investment: How are Lottery Winnings Spent?

Based on our model, for large lottery winners we expect to see investment in indivisible assets, that ultimately lead to higher income. For most small lottery winners who are well below the indivisible asset threshold, we may expect transitory impacts on savings or other divisible investments that perhaps mirror savings, such as inventory and small livestock, while they smooth their consumption of small lottery winnings over time or make divisible purchases enabled by the small grant, before returning to a poverty trap.[25]

In Table 4, we begin by evaluating differences in assets between winners and losers in order to better understand how lottery winnings were spent. We show the dynamics of these impacts by measuring treatment effects separately in the short run (first endline, 4 months after grant receipt), medium run (second endline, 18 months after grant receipt) and long run (third endline, 6 years following grant receipt). We categorize investments into consumptive assets, productive assets that are relatively divisible, and productive assets that are relatively indivisible as follows:

[24] In cases for which pre-intervention levels of the outcome were collected at midline rather than baseline, we instead control for the midline level of the outcome (where the midline also took place prior to the intervention). Specifically, business inventory, (non-inventory) durable business assets, and thus the aggregated measures of divisible and indivisible assets, are constructed at midline rather than baseline, since inventory and business durables were not disaggregated at baseline.

[25] Small lottery winners who are just below the threshold may more closely resemble large lottery winners if they are able to combine small lottery winnings with existing savings to invest in a large asset. Note that these regions are small in the upper-left panel of Figure 2 and the state is itself quite transient as households accumulate savings rapidly in this region to soon invest, so these compose a small minority of households.

- *consumptive assets:* (1) home durables (e.g., a radio or couch);
- *divisible assets:* (2) small livestock and agricultural tools (e.g., a pig or a spade), (3) business inventory (e.g., items for retail such as soap or salt);
- *indivisible assets:* (4) large livestock (e.g., cattle) and ploughs,[26] (5) business durables (e.g., a sewing machine or mirror for a salon) and (6) land.[27]

Appendix F includes descriptive statistics summarizing the purchases that respondents reported making with grants.

Focusing first on the large grant (in Table 4), we find large, positive, but ultimately transitory impacts of winning the large grant on land values. Indeed, average impacts on land values are larger than the large grant itself, which we will discuss in greater detail in Section 4.3.1. Large lottery winners also realize a transitory impact on divisible investment, perhaps using these as a temporary store of value while determining the final use of their winnings. Our primary interest is indivisible investment, so we begin with those headline results. Within 4 months of winning the large lottery, large lottery winners have already started to differentially invest in indivisible assets. At first endline, the full impact of winning the large lottery (i.e, $\beta_1 + \beta_2$) on the stock of indivisible assets is statistically significant ($p < 0.05$) with a magnitude of roughly 4 million UGX, or about $1100. At 18 months, the impact remains statistically significant and has grown to nearly 8 million UGX ($p < 0.05$). The increase is largely explained by an impact on land values, which is of a similar order of magnitude as the overall impact on indivisible assets and statistically significant.

Focusing next on the small grant, we find no statistically significant effect on total indivisible investment for the small grant winners in any of the endlines. Instead, the impacts of the small lottery show up (temporarily) in divisible investments, as predicted by our model. Total divisible investment initially increases by 241,000 UGX ($p < 0.01$), or about $67, driven by an 86,000 UGX increase ($p < 0.01$) in small livestock and agricultural tools and 155,000 UGX increase ($p < 0.01$) in business inventory. Looking next at indivisible assets, we do find the small lottery winners realize a statistically significant increase in agricultural assets (large livestock and ploughs) of 125,000 UGX ($p < 0.05$), or about $35, at first endline. This is notably much smaller than the impact of the large lottery on indivisible assets. Eighteen months after grant receipt, this point estimate has fallen but remains statistically significant, but by six years the point estimate is no longer statistically distinguishable from zero. The sum of the impacts on small livestock/tools, large livestock/ploughs, and business inventory roughly equals the size of the small grant, indicating that the small grant was largely invested.[28]

[26] The plough is the only large agricultural machine reported by anyone in our sample.

[27] In principle, land is divisible. In practice, we rarely see purchases of land below 500,000 UGX — only about 5% of land transactions made recently, in June 2013 or after, are valued at less than 500,000 UGX — suggesting that land is typically sold in larger units. Small, disbursed, parcels of land are unlikely to be as productive as large, connected plots; as a result potential land buyers must either restrict themselves to purchasing from neighbors or purchasing a large plot on its own. The matching process between buyers and sellers can create inefficiencies due to the thinness of particular land markets as shown by Bryan, de Quidt, Silva-Vargas, Wilkening, and Yadav (2022).

[28] Note that regression coefficients between sub-categories need not perfectly add-up for two reasons: (1) in each regression we control for prior outcome levels, so control variables change between specifications, and (2) we winsorize each outcome separately.

Comparing the impact of the large versus small lottery on divisible assets, we see that the full impact of winning the large lottery (again, $\beta_1 + \beta_2$) on divisible investments is 256,000 UGX, or $71, which is marginally significant ($p < 0.10$) and includes a statistically significant increase in business inventory of roughly 310,000 UGX ($p < 0.01$), or $86. This is dwarfed by the 4,000,000 UGX ($p < 0.05$), or $1100, impact on indivisible investment and constitutes a smaller share of total winnings. Indeed, we cannot reject that large and small lottery winners invest the same amount in divisible assets (large lottery winners invest only 16,000 UGX, or 7% more, in divisible assets than small lottery winners — and this differential impact is statistically indistinguishable from zero — despite their winnings being five times as large). While any empirical comparison between the impacts of small and large lotteries is not experimental (the recipients differ by their own self-selection) it is still consistent with the theory. Finally, although home durables (for example, a radio, television, refrigerator, or carpet) may be one area where we would expect grant winners to spend money, especially those with no business opportunities, we find no statistically significant effect on home durables for either the small lottery or large lottery winners at any of the endlines.

The impacts of both types of grants on assets are ultimately transitory. While the impacts of the large grants on indivisible assets persist and even grow at 18 months, by 6 years there is no longer any statistically discernible impacts. The impacts on divisible investments are even more transient, no longer significant at 18 months, except for a marginally significant ($p < 0.10$) impact of the small grant on business inventories, but the point estimate is less than half its value at 4 months. These large but temporary increases in divisible assets are consistent with their use as a temporary store of value for the winnings. In sum, neither the impacts of small or large grants — including the initially large impacts on land values for large lottery winners — are sustained over the long term. In the following subsection we further investigate the transitory, large impacts on land values reported by large lottery winners.

Relating back to the model, we emphasize three important pieces of information. First, those who chose and then won the large lottery invested in indivisible assets. Second, the dominant indivisible asset is land. And, third, the impacts on assets have diminished substantially by third endline. The first is consistent with the theory of poverty traps, but the latter two will inform our aggregate results that financial services are less powerful in addressing poverty traps in aggregate and that poverty traps are not as permanent as the term might suggest.

4.3.1 Unpacking the Large but Transitory Impact on Land Values

Our headline results show a temporary increase in land values for large lottery winners, including the surprising facts that (i) the impact of the large grant on land values at 18 months far exceeds the size of the grant itself and (ii) has dissipated six years later. Both facts merit further scrutiny. Treatment effects on land values for lottery winners should only come from investments that they made in the land: either to

purchase more land or to invest in land improvements, maintenance, or building.[29] Beyond the scope of our parsimonious model, there are several potential explanations for these transient, large treatment effects on land values:

1. A temporary increase in land demand for the first two years following the grant caused in part by the grants we offered.

2. Misallocation prior to treatment: lottery winners are able to invest in land that improves their productivity. This land may have a high private value because of the impact on their productivity, but would not have any impact on market land values, which may be revised down by the participant over time.

3. A temporary speculative land value bubble of which lottery winners are better informed (because they are more likely to be active in the market for land).

4. Investments in buildings and assets on the land, which depreciate over time.

5. Initial large land purchases as a temporary store of value, which have been sold by the third endline (at 6 years).

6. Unmeasured transfers from others that enable purchases or investment on existing land and exceed the immediate household's available resources (or unmeasured transfers from own savings or own assets that enable purchases of new land or investment on existing land, also exceeding the size of the grant), which are ultimately returned.

These explanations differ in their implications for the long-term impacts of cash grants or expansion of financial services. For example, grant winners making already owned plots of land more productive has distinct implications for aggregate growth relative to grant winners simply speculating in land and driving up its price, especially if land is in fixed supply. If land valuation reflects market driven shocks, whether aggregate or idiosyncratic, this would also have important implications, as we will see in Section 5. Moreover, if demand for land is driven by returns to speculation this would have repercussions on the external validity of our finding for the importance of land. However, land appreciation in semi-urban areas is quite common — for example, as Korah, Matthews, and Tomerini (2019) discuss in the areas surrounding secondary cities in Tanzania and Ghana.[30]

To confirm whether our measured impacts on land investment reflect actual expenditures, we examine the extent to which large grant winners made new land purchases and investments in improving land (both of which could increase a household's land value) first through descriptive statistics in which we simply asked respondents directly how they spent the grant in Table F.1. About one-third of large lottery winners

[29] There may be real appreciation in land values that increases the value of land for all land owners, across both treatment and control groups, and, indeed, as noted above we find substantial capital gains in land over our study period (see Appendix A). These capital gains should not differentially impact treatment or control group land values, however.

[30] Moreover, the idea that land values in urbanizing regions appreciate at a rate faster than can be accounted for by internalizing the future flow of income from the land has a long history, including, e.g., Clonts (1970); Deininger et al. (2003).

report that a land purchase accounts for the largest expenditure made with grant winnings; a land purchase is also the modal response when large grant winners are asked what accounts for the biggest purchase made with grant winnings. About 6% of small grant winners also report that purchasing land was their largest expenditure with the grant.[31] 14% of large grant winners and 9% of small grant winners report that land and building improvements accounted for the largest fraction of their grant spending, which would have increased the value of their land as well. We conclude that land values likely reflect a combination of new land purchases and land improvements, with land purchases being the most common expenditure incurred with the large grant.

To further differentiate whether lottery winners purchased land for investment versus speculation, we ask respondents why they hold land and find that large lottery winners are indeed more likely to hold land for income generation and less likely to hold land for speculation. Appendix Table F.2 shows that large lottery winners are 9.5 percentage points more likely to hold land for income generation than small lottery winners ($p < 0.10$) and 12 percentage points more likely to hold land for income generation than the control group (not significant, p-value=0.12). By contrast, large lottery winners are statistically significantly less likely than small lottery winners to hold land for the purposes of speculation or social reasons (such as passing inheritances to their children), and less likely than the control group in general to hold land for speculation ($p < 0.05$).

4.3.2 Land Transactions and Use

In our third endline (6 year) survey, we asked respondents retrospective questions about land purchases, investment, and sales. In Table 5, we show that winning the large lottery differentially increases the likelihood of purchasing land between the midline and third endline by 15 percentage points ($p < 0.05$) relative to small lottery winners and by 16 percentage points ($p < 0.05$), or 48%, relative to the large lottery control. Notably, large lottery winners are no more likely than either small lottery winners or the control group to have sold land or made investments on existing land between midline and the third endline.

Large lottery are both more likely to purchase land and to make purchases that add land to existing plots. Large lottery winners are 9.2 percentage points more likely than small lottery winners (approaching marginal significance) and 6 percentage points (not statistically significant), or 50%, more likely than the large lottery control to have added land to an existing plot between 2016 and the third endline in 2023. Large lottery winners tend to add land to their *least* cultivated plot: they are 5 percentage points ($p < 0.05$) more likely than small lottery winners and over twice as likely as the control group (not statistically significant) to add land to their least cultivated plot between 2016 and the third endline. In contrast, large (and small) lottery winners are no more likely than the control to add land to their most cultivated plot. This is important if lottery winners are able to make formerly less productive plots more productive by increasing their size.

[31] A few small grant winners purchasing land is also consistent with the model's predictions that some households whose baseline wealth is sufficiently close to the threshold above which they can make the indivisible investment will choose the small lottery, if winning the small lottery provides enough funds to surpass the threshold.

After examining the impact of the grants on income, we will revisit the returns to adding land to existing plots in Section 4.4.4.

That large lottery winners are both more likely to purchase land and that a meaningful share of these acquisitions consist of adding to existing plots is consistent with newly purchased land having a higher private value to the lottery winners than its market price would reflect. Moreover, we find no impact on the number of plots or on the average size of a plot, so these purchases were likely small quantities of contiguous land — consistent with having high private value but not necessarily greater market value than other land.

4.4 Other Impacts on the Household Budget: Consumption, Income, and Savings

We have found substantial appetite for a riskier, larger potential payout lottery and subsequent investments in large indivisible assets among those who chose the riskier lottery and won large grants, consistent with the poverty trap model. Using equation 7, we now consider the impact of the lottery on other (non-investment) components of the household's budget constraint: consumption, income, savings, and net credit. Consumption and income are of interest, since the model predicts that if large grants allowed those choosing the riskier lottery to escape poverty traps through indivisible, high yield investments, then these higher yields should manifest in increased income and ultimately consumption. Savings and net credit are of interest in understanding the overall cash flow and budget of households.

Table 6 presents the results. Again, we focus first on total impacts for large lottery winners ($\beta_1 + \beta_2$). We indeed find increases in both consumption and income. As in the model, increases in consumption take time to materialize. Large lottery winners do not have higher consumption at the first endline, but do by the second endline, 18 months after receiving lottery winnings. Large lottery winners realize 44,000 UGX ($p < 0.05$), or \$12, higher consumption per month, which is 18% greater than the control group. Six years after grant receipt, the impact on consumption falls only slightly to 40,000 UGX, though its statistical significance has dissipated ($p = 0.17$). The fall in statistical significance reflects a larger variance, consistent with shocks hitting households over time, as in the model. We also find evidence of gradual income gains that are broadly consistent with the model, but the timing is more complicated. Large lottery winners' monthly income does not appear to change until the third endline, when it is 90,000 UGX ($p < 0.05$), or \$25, higher per month. The lack of increased income at second endline is perhaps puzzling given the increase in consumption, and we will return to this.

We note the impact on savings and credit. We find a statistically significant ($p < 0.05$) increase in savings of 160,000 UGX, or \$44, at the first endline, but this slowly falls in later endlines, perhaps marginally significant ($p = 0.11$) at 18 months and completely insignificant by six years. Point estimates on net credit are negative in the first two endlines but never statistically significant, indicating that the grants were neither used to repay debt nor to leverage loans for indivisible purchases. This is not surprising given the overall low levels of financial services in the sample at midline.

Focusing next on the small lottery winners, we find no significant impact on either consumption or income from winning the small lottery at any point, though point estimates are positive at first and second endline. Similarly, we find no significant impacts on borrowing, though point estimates are negative. Savings increases at first endline by 72,000 UGX ($p < 0.05$), or \$20, but statistical significance is lost by 18 months and the point estimates dissipate over time. Taken together with the impacts on divisible assets, this is consistent with small grants being used to save in various forms in the short-run but being spent down and ultimately dissipating in the longer run.

Is it surprising that there was a relatively large income impact of the large lottery at third endline and consumption impacts at both second and third endline? To assess this, we try to better investigate the distribution and source of income gains, showing that income gains (i) concentrate large lottery winners into the upper tier of the income distribution, (ii) are primarily agricultural, (iii) partly reflect exogenous agricultural price shocks, and (iv) partly reflect productivity gains.

4.4.1 Distribution of Income Gains

In Figure 4, we use distribution regressions (e.g., Chernozhukov, Fernández-Val, and Melly (2013); Goodman-Bacon (2021)) to show how the treatment effect emerges over the income distribution at each endline, for both small and large lottery winners. The vertical dotted lines correspond to the 25^{th}, 50^{th}, and 75^{th} percentiles. At the first endline, large lottery winners are not statistically significantly more likely to have income in the higher (or lower) part of the distribution than the large lottery control group. By the second endline, large lottery winners are about 5 percentage points more likely than the control to have monthly income above 380,000 UGX, or \$106 (corresponding to the 75^{th} percentile), but this is imprecisely estimated. Large lottery winners may also be less likely to have income in the middle of the distribution — perhaps indicating that they are now more likely to compose the upper end of the income distribution — but again, these effects on income are still imprecise at the second endline. By the third endline, however, large lottery winners are statistically significantly more likely (specifically, about 12 percentage points more likely) to have monthly income that exceeds 600,000 UGX, or \$167 (corresponding to the 80^{th} percentile of the distribution). This suggests that large lottery winners disproportionately compose the upper end of the income distribution, as a result of receiving the large grant. While not statistically significant until about 600,000 UGX, the effect of the large grant on the probability of having income greater than X at the third endline is positive for all values of X that exceed 260,000 UGX, or \$72, which is just over median income. In other words, the large lottery moved its recipients into the upper half of the income distribution, and the distribution regressions at each endline underscore that this effect emerges over time and only for the large lottery winners, consistent with households in the virtuous cycle of an economy with poverty trap dynamics. At the third endline, small lottery winners do not disproportionately compose any part of the income distribution relative to the control, consistent with any gains yielded by the small lottery being quite transient.

Although we find higher income, we do not detect statistically significantly higher investment levels persisting through the third endline. The higher income is *not* a temporary effect of immediate income from having just sold assets, but there are several potential explanations to reconcile higher income without higher asset values, and indeed the explanations are no doubt varied across households.[32] To highlight the primary mechanisms, we now examine the source of income gains.

4.4.2 Decomposing the Impact on Income

We evaluate the extent to which the income gains can be linked directly to indivisible investments made with a large lump sum of capital, as the theory suggests. We begin by decomposing the source of the large lottery impact on total household income, by estimating the extent to which it comes from business income, crop income, livestock income, salary income, or remittances. For this discussion, we normalize coefficients by the impact on total income, so that each coefficient is a percentage of the large grant's impact on total income (and the sum of the impact of each sub-component is 1). Standard errors are calculated using the delta method.

Results are shown in Table 7. Recall that winning the large lottery led primarily to investment in land, alongside short-run increases in business inventories and indivisible agricultural investments (e.g., large livestock). Consistent with this, increases in crop income are statistically significant ($p < 0.05$) and explain roughly half (49%) of the total increase in household income. Growth in business income explains another quarter (27%), and livestock income, salaries, and remittances explain the remainder, but only crop income is statistically distinguishable from zero.

4.4.3 Crop Prices

Our theory emphasizes the importance of shocks for poverty trap dynamics, but we abstract from aggregate shocks. Such shocks clearly exist in the real world, however, and can confound both external validity and interpretation of RCT results (Rosenzweig and Udry, 2020). Through the lens of the model, neutral aggregate shocks would not impact treatment relative to control but an aggregate shock to entrepreneurial productivity would disproportionately impact those who have high entrepreneurial income via higher productivity and/or more capital. Using crop income and land as the empirical counterparts, we explore whether aggregate shocks to the global price of crops could explain some of our income treatment effect.

We construct an annual index of the prices of crops relative to overall consumption. In our sample, the three most common crops are bananas ("matooke"), beans, and maize. Price data are only consistently available for the latter two and only for neighboring Rwanda, but are presumably correlated with commodity

[32] Many of the durable investments measured at 18 months may continue to generate income for the participants, but they are likely to have depreciated significantly over 6 years. While the market value of the durable depreciated assets is quite low, the assets may continue to be productive assets in the business. Improvements made to land may have similarly depreciated, increasing long-term productivity while having little impact on overall land values. Temporary improvements made to businesses may have also yielded long-term clients who remain customers and have increased profitability. We note that depreciated capital (whether physical capital or "customer capital") without further investment should eventually lead to income reversion, with income following productive assets. In Section 5, we demonstrate reversion of both in the aggregate shares.

prices in Uganda for globally traded grains.[33] We therefore use these data to construct relative price indexes for Uganda.

Two things are of note in the series. First, relative prices of beans and maize in 2023 were particularly high, 76% and 123% over their mean relative price during the sample period of 2015-2023, and such prices can reasonably be taken as exogenous to the Ugandan farmers in our sample. A simple back-of-the-envelope calculation shows the quantitative importance of this. With negligible expenditure on inputs, a 100% gain in price (the midpoint of the price increase in beans and maize) would result in a 100% gain in crop income. With crop income constituting 30% of total income at midline for those selecting the large lottery (Table 3), this would lead to a 30% increase in income overall, a seemingly sizable impact. The treatment effect on income at third endline is 90,000 UGX, or a 24%(=90,000 UGX / 373,000 UGX, as in Table 6) increase *relative to the control group*, with an increase in crop income accounting for half (in Table 7), or 12 percentage points, of the total treatment effect on income. For an increase in crop prices to explain the treatment effect on income, crop output must be disproportionately higher among treated households. Thus, quantitatively, for an increase in crop prices to fully explain the income increase, the output of treated households would have to be 40% (i.e., $\frac{12\%+30\%}{30\%}$ - 1) greater than the control.[34] Given comparable inherent productivity (ensured by the randomization), this would presumably reflect higher effective agricultural inputs (k in the model), particularly land or effective land, which we will examine in the next section.[35]

Second, we note that crop prices for beans and maize in 2018 were 25% and 50% below their mean, respectively. Thus, underlying output differences may well have been sizable (and consistent with the model) despite our not observing an income difference in the second endline. Low prices may also have led agricultural households to disproportionately consume more in that year.

4.4.4 Improvements in Productivity or Effective Land?

As noted, crop prices may drive income effects, but differential crop income gains across treatment and control require differentially higher crop output, presumably from greater land or effective land. We measure higher land values at second endline but not third endline. We now examine whether "effective" land might be higher because it is used more productively. This idea of effective land goes outside our parsimonious model, but Foster and Rosenzweig (2022) show that there may not be returns to scale until farmers reach a level that can fully utilize gains from mechanization, and these farmers, with an average

[33] The FAO reports annual prices for crops but not for Uganda. The closest available country is Rwanda, for which data are available annually for common beans and maize. (Data for bananas are only available for 2015 and 2020.) We take Rwandan price indexes in USD, convert them to UGX using the official exchange rate, and construct relative price indexes by dividing by the Uganda CPI. Kenyan price indexes are also available, but the USD series are not complete, importantly lacking data for 2017 and 2018. Across available common years, 2015-2023, the correlations with Rwanda are high though not perfect: 0.66 and 0.60, for beans and maize respectively. The Kenya data also show sizable upward price shocks in 2023. Importantly, these price effects are likely over-estimates of local effects as these periurban areas are likely neither importers nor exporters of agricultural commodities. More likely, any productivity impact on the treatment group would *reduce* local prices for both the treatment and control group.

[34] In our data, the ratio of crop revenue to crop income (net of the costs of inputs like seeds, fertilizer, and hired labor) is 134%, or roughly four-thirds. Accounting for this, a 100% increase in crop prices would double revenue but increase income by 134%–assuming no increase in input prices. Hence, accounting for input costs, we would only require output to be 30% more than the control.

[35] We found no significant impacts on other input usage.

of 1.2 acres, are below that level. On the other hand, productivity may also improve from having more contiguous land and land near the household — a possibility that we address given the result, in Table 5, that large lottery winners are differentially likely to add land to existing plots.

In Table 8, we show that adding land to an existing plot is associated with higher crop income, conditional on total acres cultivated and the number of plots. First, we see across multiple specifications that conditional on total acreage, splitting cultivation into separate plots is associated with a reduction in crop income. While total acreage is significantly positively correlated with crop income, the number of plots is significantly negatively correlated with crop income. Adding land to an *existing* plot, however, is associated with an increase in crop income of 26,000 UGX ($p < 0.05$), or \$7, per month. This is equivalent to a 26% or a 0.20 standard deviation increase in crop income. Adding land to the least cultivated plot leads to a similar 27,000 UGX increase in crop income, though this association is not statistically significant. Together with the earlier results on land transactions from Table 5, these results suggest that (i) large lottery winners expanded cultivation by adding land to plots they already farmed rather than acquiring entirely new independent plots, and (ii) that such within-plot expansion is associated with meaningful gains in crop income. This is consistent with the private value of adjacent land to existing plot-holders exceeding its open-market price — a pattern not captured in reported land valuations but reflected in the productivity gains observed over the long run.

5 Calibrated Model Results

We now turn to our analysis of the macroeconomic implications of indivisible investments using our model of Section 2. To calibrate the model, we use the income and savings distributions observed in our midline data and the households' choice of lottery and the pre-treatment transition probabilities of income and savings positions to calibrate the model. It is common in the macro development literature that links experimental data with macro models to use the experimental data to either calibrate or cross-validate the model. We do the former, using the observed, experimental changes in income, capital, and liquid assets at the various endlines to calibrate the model. Having done so, we use the model to quantitatively assess the long-run impacts of cash grants and financial intermediation on development.

5.1 Last Choices for Aggregate Model

We start by closing the model of a household developed in Section 2. If land is elastically supplied, we take its price, P, as given, and aggregate using expectations over the distributions. When land is in fixed supply, its price is determined by the market-clearing condition

$$\int k_{it+1} \, \mathrm{d}i = \bar{K}$$

determining the equilibrium price. We normalize units so that $P = 1$ in the initial steady state prior to interventions.

Mapping the model to data requires choosing functional forms to match heterogeneity in the data. Permanent heterogeneity in labor income, entrepreneurial productivity, as well as the discrete shock on the ability to operate land allow the model to match the observed distribution of land holdings, income and liquid assets in the data, as well as reproduce the responses to our lottery experiments. We choose independent log normal distribution for $\bar{e}_i$, the permanent component of labor income, and $\bar{z}_i$, the permanent component of entrepreneurial productivity, which are drawn at birth from $N(\log \bar{e}, \nu_e^2)$ and $N(\log \bar{z}, \nu_z^2)$, respectively.

Mapping to lottery choices merits acknowledging that the lottery choices we observed may reflect other factors that induce risk-loving behavior (simple mistakes or misunderstanding, for example). Our analysis in Section 4.1 gives us confidence that they are not purely random, but there may still be noise. We therefore allow for preference shocks that may lead to risk-loving choices even away from the investment thresholds. Specifically, we first convert the consumption equivalent value functions back to utils, so

$$\hat{V}^j = \frac{1}{1-\beta} \log V^j$$

and suppose the agent derives utility $\hat{V}^j + \varepsilon^j$ when choosing lottery option j, where ε^j is drawn independently from a Gumbel distribution with scale σ. This yields the standard multinomial logit choice probabilities

$$\Pr(\text{choose } j) = \frac{\exp\left(\hat{V}^j/\sigma\right)}{\sum_{\tilde{j}} \exp\left(\hat{V}^{\tilde{j}}/\sigma\right)}.$$

When $\sigma \to 0$ the model collapses to the deterministic choice described above; larger σ smooths choice probabilities toward a uniform distribution. The orthogonality of logit specification in utility allows for an interpretation of the taste shocks as simple measurement error in the lotteries that does not interact with model dynamics. We estimate σ jointly with the other structural parameters.

Lastly, we need to choose a functional form for the discrete levels of capital. We choose quadratic spacing following the formula:

$$k_j = k_{min} + (k_{max} - k_{min}) * ((j-1)/(n-1))^2, j = 1, ..., n$$

with $n = 4$ for a total of five grid points.

5.2 Calibration

The calibration then proceeds as follows. We start by assigning four parameters. First, we set the maximum loan-to-value ratio $\theta = 0$ in the benchmark. We do so for simplicity, but this broadly matches the limited access to credit that we observe. The interest rate r then effectively becomes the interest on savings,

which we set to zero, consistent with the fact that our empirical sample all were provided with zero interest savings accounts. A period is a month, and we set the monthly maintenance cost of capital, δ, which is analogous to depreciation in an indivisible capital model, to 0.01. We set the constant survival probability, p, to $1 - 1/480$, implying an average adult life (or career) of 40 years.

We are left with ten parameters to calibrate: the discount factor, β; the parameters disciplining the capital technology, k_{min} and k_{max} (the minimum scale and maximum levels of indivisible capital, respectively), those disciplining the entrepreneurial productivity, $\bar{z}$, ν_z, and σ_z (governing the (log) mean, permanent dispersion, and shock dispersion, respectively); the parameters disciplining the labor income: ν_e and σ_e (disciplining the dispersion of the persistent income difference and productivity shocks, respectively), and the states disciplining the capital-ability shocks, i.e., the persistence of the non-entrepreneurial state ($\xi = 0$), p_L, and the entrepreneurial state $\xi = 1$), p_H.

Our calibration strategy is to choose these parameters to match the moments in the data that are most of interest to this study: the observed lottery choices, the distribution and dynamics of income and assets, and the experimental treatment dynamics. For the lotteries, we target the percentage of households who choose the risky lottery. For the distributions of assets, we target both land, and liquid assets separately, and for each of the three distributions (income, land, and liquid assets), we target the mean as well as five percentiles of the distribution (10th, 25th, 50th, 75th, and 90th), a total of 14 moments given the normalization of median income to one in the model and in the data.[36] For the dynamics of income and assets, we target transitions across income terciles over the 16 months from baseline to midline. For income, we report transition probabilities over four months, while for land we scale them to measure the transition probability over 58 months. Each column of two 3x3 transition matrix constitutes sum to 1, so each matrix constitutes only six independent moments, for a total of 12 additional moments.

Finally, we are interested in aggregate predictions. Rather than targeting the micro treatment effects from regressions, we target the macro distribution of income, capital, and liquid assets across the four groups (large lottery winners, large lottery non-winners, small lottery winners, and small lottery non-winners). (We use distributions rather than aggregates because the different survey waves may reflect aggregate shocks independent of the experiment.) We target baseline numbers as well as the changes between baseline and the first two endlines. Again, for each moment, the four lottery groups sum to one, so we have three independent lottery groups, three outcomes, and three survey waves yield a total of 27 additional moments. Together, these targets amount to (1+14+12+27) 54 moments, substantially more than the ten free parameters we calibrate.[37] We retain 9 independent distributional moments from the third endline for out-of-sample evaluation.

Given the parsimony of the model, we cannot match all moments, and instead minimize a weighted

[36] Savings includes only financial wealth and corresponds to S in the model.

[37] We examined alternative models with additional parameters such as non-unitary risk aversion, Epstein-Zin preferences (with separate risk aversion and intertemporal substitution parameters), returns to scale in land, both physical and pecuniary fixed costs to adjusting capital, but these extensions did not significantly improve the fit over our parsimonious model.

average of the sum of percentage squared deviations.[38] Table 9 summarizes the parameter values and model fit.

Focusing first on Panel B, we note a few patterns. First, recalling that our model is a monthly model, the discount factor is low (but not inconsistent with the fact that few chose to delay their winnings one-month at a 3% interest rate). Uncertainty is significant. The dispersion parameters of the shocks to both labor productivity and entrepreneurial productivity exceed those of the permanent cross-sectional variation. The non-entrepreneurial shock is also important, but the persistence of the entrepreneurial state is higher than that of the non-entrepreneurial states, implying that an expected spell in the former is about 6 years, while the latter is less than two years. Correspondingly, households spend 77% of their time in the entrepreneurial state. The indivisibility is also sizable, with the minimum scale of a land purchase being over six months of median income and four times the median holdings of liquid assets. The maximum scale of land is 20 times higher than the minimum scale, allowing for a strong concentration of land among the wealthy tail.

Panel C examines the model fit in three parts. Despite the extreme over-identification, the model fit is reasonably close, especially for our purposes. In C.1, although the model considerably overstates average income and liquid assets, the distributions are reasonable, showing leftward skews, means above medians, a high concentration of income in the upper tail, and an even higher concentration of liquid assets. For land, we also have a strongly leftward-skewed distribution that has concentration of land in the upper tail, and a landless population as well, but we grossly understate land levels overall across the distribution. Given the ε^j term, the model is able to exactly match the share choosing the risky lottery, but the parameter explains only 28.5% of risk loving choices. The remaining nearly three-quarters of risk-loving choices reflect otherwise optimal risk-loving behavior induced by the indivisibility of investment.

Panel C.2 reports the distributions of income and assets across the different lottery subpopulations and how they change over time in response to the lotteries. (Note that although we report changes for controls, SUTVA is not assumed to be violated, since controls change only as a share of the total.) For comparison, the top row reports the prevalence of the group in the population. Focusing on those who chose the large lottery, our primary focus with respect to the consideration of indivisibilites and poverty traps, we see that those who chose the large lottery are richer on average at baseline. The model captures this feature. At first endline, the winners and losers of the large lottery have diverged in terms of income and land, and the model captures this. In the model, the shares of large winners have changed by 0.5 and 2.0 percentage points, respectively, of comparable magnitudes to the 1.1 and 1.5 percentage points in the data. The impact on land stays relatively stable through the second endline while the share of income increases, and this is true in both model and data, although the income increase is stronger in the data.

Again, the results for third endline are not targeted but we report them as out-of-sample moments. The model again does reasonably well: By third endline, both the increases in the shares of income and capital accounted for by large lottery winners have dissipated from their peaks in the second endline. Indeed and

[38] The weights are equal to one, except for the lottery choice moments, which are critical in getting the indivisibility dynamics and so get a higher weight of 20.

of note: in both model and data, with the exception of liquid assets among non-winners in the data, the magnitudes of all the numbers overall have declined. This is an important feature of both model and data that will speak to the long-run impacts of cash grants.

Panel C.3 shows the tercile transition matrices for income and land for both model and data. In the data, the diagonal is strong, with the middle tercile showing the least persistence, and the the diagonal being stronger for assets than income. Nevertheless, the off diagonal elements are not negligible, with the lowest number being an 8.2% chance of transitioning from low income to high income four months later, and a 9.3% chance of transitioning from low wealth to high wealth 48 months later. The empirical matrix shows a slight tendency to polarize, but it is very slight, implying an ergodic distribution with 39% and 35% in the lower and upper terciles but 30% in the middle tercile. This churning is an important feature of the data that belies poverty traps in the sense of permanent vicious (or virtuous) cycles. Instead, given the poverty trap dynamics that prevail for a given productivity state, the model calibrates considerable shocks in order to match the high level of churn, and indeed both transition matrices are reasonably well matched by the model. Again, this high level of churn speaks to the long-run impacts of cash grants.

5.3 Assessment of Policies for Long Run Development

We now use the model to evaluate the development impacts of two common policies suggested to address the presence of poverty traps: cash grants and financial services. Macro development models are useful for thinking about long run dynamics and general equilibrium. We emphasize the long run dynamics in the first exercise, and the importance of general equilibrium in the second. Within the model, the mechanisms for promoting development are two fold: 1) increasing the amount of capital used, and 2) improving the allocation of capital across producers with heterogeneous productivity, i.e., reducing misallocation in the sense of Restuccia and Rogerson (2008); Hsieh and Klenow (2009).

To assess cash grants, we simply extend our model's predictions for the treated groups out beyond 10 years after treatment. We stick with this design as approximating a best case scenario in that the grants are (i) targeted to those who are in the convex region that characterizes the poverty trap, and (ii) small scale so that general equilibrium effects do not undermine the impacts. Finally, our model has death, but we plot treatment effects only for survivors, again to make a conservative analysis. Regarding targeting, although our experimental data lack the cross-diagonals of treatment, i.e., people playing the lottery they did not prefer, we can easily simulate these counterfactuals in the model to see how targeting impacts treatment. We return to general equilibrium impacts in our later counterfactual policies.

Figure 5 shows treatment effects on income and capital over time. To measure treatment effects, we calculate the average change from midline (indicated by the vertical black dashed lines at month 25) as a percent of the midline level for both winners and loser of a given lottery, and then difference the two. The solid blue lines are for the actual (preferred) lottery chosen and correspond to the empirical experiment and the targeted policy. The dashed red lines illustrate the hypothetical treatment effects if people had been

granted the lottery they did not prefer. The left column of panels shows the treatment effects for winning the large lottery, while the right shows the small lotteries, with the top panels showing income and the bottom showing capital.

Focusing on the large grant panels on the left and actual grants (solid lines), the peaks are substantial, constituting an over 25% increase in income and a four-fold increase in capital, with each coming roughly six months after winning. The impacts are persistent but not permanent, even in this best-case type of scenario. The half-life or the properly targeted large grant is roughly 9 years for income and about 8 years for assets. Looking at the dashed lines, the hypothetical impacts of winning the large lottery are substantial for those who didn't choose it. Nevertheless, the impacts are less than one-third of the targeted impacts, and much of this comes from the fact that targeting is imperfect, given the preference shock for the lottery. Still, the later peak indicates that the grants do induce even those without an indivisibility-driven preference for risk to invest in capital, but it takes them longer to save enough up. In sum, the calibrated model predicts fade out of both assets and income over time, just as we found in the empirics.Given the lack of aggregate shocks in the model, as observed empirically in prices in our third endline, the fade away in income in the model is somewhat quicker than in the data. We discuss aggregate shocks below.

Looking at the right columns, we see that the small grants have much smaller treatment effects, but even these treatment effects are persistent, with half lives of just under six years for income and five years for assets. The peaks for the properly targeted grants are nearly immediate, but delayed for the hypothetical improperly targeted grants, indicating that the small grants do help some risk-loving agents to eventually invest in indivisible capital.

Given the poverty trap set up, targeting is key to sizable impacts. Comparing across the figures, the impacts on income of the large grants are over ten times those of the small grants, which is much more than the under five-fold difference in the size of the grants. Much of this difference comes from targeting, however, as there is less than a two-fold difference in the peak levels between the large and small grants, when improperly targeted. Again, the impacts are persistent but not permanent. That is, the calibrated model predicts fade out of both assets and income over time, just as we found in the empirics.[39] The reason is that shocks hit the economy, and there is mean reversion in states and corresponding mean reversion in economic outcomes. Again, the several-fold features of the data that support this prediction are both experimental and observational: (i) the disappearance of the treatment effect for land at third endline, (ii) the declining share of land and income at third endline relative to second endline, (iii) the substantial economic movement in the transition matrices for land and income (despite high levels of inequality), and (iv) the resulting high variance of shocks required for the model to replicate the data. One-time cash grants can have highly persistent impacts given the presence of poverty trap household dynamics, akin to the results in Buera and Shin (2013), but despite these dynamics, one-time cash grants do not appear to be a policy for long-run development in these communities.

[39] Given the lack of aggregate shocks in the model, as observed empirically in prices in our third endline, the fade away in income in the model is somewhat quicker than in the data. We discuss aggregate shocks below.

A less costly, more sustainable alternative to cash grants may be introducing wide-scale financial services to the regions of study. We assess four different counterfactuals relative to our benchmark.

As emphasized in the introduction, we show that the elasticity of supply of capital is crucial here. For each counterfactual, we therefore assess these policies under two different scenarios about the elasticity of the supply of capital. The first is a scenario where capital is perfectly elastically supplied at a constant price, as is typically assumed in macro models of investment. In this case, only quantities will adjust to an increase in the demand for investment. Such an assumption is surely appropriate when the intervention is small relative to overall supply of capital, and so we label this "PE". The second scenario is one in which capital is in fixed supply. Such a scenario is an extreme scenario but perhaps more appropriate for the case when land is the investment good. In this case, prices will adjust to changes in the demand for capital rather than aggregate quantities, and so we label this "GE". We note, however, the purpose of both scenarios is to use the extreme cases of the elasticity of supply to bracket possible impacts.

Table 10 presents the results for the various counterfactuals under both scenarios. We focus on the impacts of these counterfactuals on aggregates: income, consumption, net savings, capital, the value (price $\times$ quantity) of capital, the price of capital, and capital income. Given our interest in poverty traps and wealth dynamics, we also report the fraction of the poor in the economy (i.e., the fraction below a threshold defined by the lowest tercile *of wealth* in the benchmark economy) and the probability of staying poor. We report all of the counterfactual results relative to the benchmark values (which are given in the notes of Table 10 but not otherwise of particular interest). We point out that although savings and capital are positive, consumption approximately equals income in the steady state because capital does not depreciate and income is already net of maintenance costs. That is, steady state investment is zero in the economy whether capital is elastically or inelastically supplied. The only difference between consumption and income comes from the small probability of death.

The first counterfactual we evaluate is a 'Divisible Capital' case in which capital is perfectly divisible (though still bounded above) enabling a choice of any $k \in [0, k_{max}]$. This might be a viable policy counterfactual if indivisibility in land, for example, were driven by the titling of plots or rental markets frictions. (However, our qualitative research suggests that titles are *not* the driving characteristic of land indivisibility.) Alternatively, this exercise can be understood as providing a benchmark comparison for how powerful indivisibility in investment can be in driving economic aggregates and poverty in order to answer the question in the paper's title. Comparing the first and second columns, we see that indivisibility substantially lowers output. In the PE case, output would be 82% higher without the indivisibility, and capital would be over three times as high, since even the poor-but-productive would be able to invest some, and gradually expand their capital over time. Both the fraction poor (lowest income tercile), and the probability they stay poor four months later are lower in this world. The GE column shows that the aggregate impacts in PE are wholly the result of capital accumulation. In GE, when the stock of capital is fixed, divisible capital

simply leads to an increase in its price.[40] Hence, the importance of indivisibilities for development depends critically on the common assumption that capital is elastically supplied.

We now turn to three counterfactual financial policies which we compare to these two benchmarks. The first is an expansion of credit services that enables households to borrow up to (a relatively modest) 25% of the value of their capital, which we model as an increase of the collateral parameter, θ, from 0 to 0.25. The second is an expansion of interest-bearing savings facilities, which we model as an increase in the monthly interest rate, r, from 0 to 0.03 (equal to the delayed payment interest we offered, and which was chosen by 21% of the sample). The third is a combination of both, which enables borrowing together with saving but both at a positive interest rate.

The pure credit intervention is quite powerful when the price of investment is fixed. Under this scenario ('Credit' and 'PE' in the table), the 2.51 indicates that the capital stock is 151% larger, while the 1.56 indicate that income is 56% larger. The 0.89 indicates that net savings declines, a combination of both credit and decreased savings for self-financing reasons, but by only 11%. The fraction poor declines to 47% of its original level, although the chances of staying poor remain essentially the same.

In striking contrast, in the case of credit under a fixed stock of capital, we see absolutely no impact on aggregate income, capital income, or poverty. Indeed, the policy is completely neutral even on the micro level, so it provides no benefits from reducing misallocation. The change in θ is completely offset by a commensurate change in price, where $P = 1/(1-\theta)$ or 1.33 in this case. Net savings declines as households simply keep their surplus savings above available credit, i.e., $\hat{a} = a' + \theta P k'$, unchanged. Indeed, given our assumptions, one can easily show this neutrality result analytically.[41] The key point is that the power of credit to increase output and reduce poverty is substantially reduced in a world where the increased demand for investment manifests itself more in an increase in prices than through an increase in productive capital.

The results for the pure savings intervention demonstrate that when the investment price is fixed, interest-bearing savings can be an even more powerful financial intervention than credit. Under a fixed price of capital, savings responds strongly to the positive return on savings, increasing more than eight fold, a qualitative contrast to the impact of the credit policy. Similarly, capital increases nearly four fold, capital income increases more than three fold, and overall income is a near tripling. The impact of savings on

[40] Recall, that our production has constant returns to scale. An earlier version considered diminishing returns to scale, the common assumption for capital, which leads to an increase in output from the existence of an Inada condition at the now possible low levels of land. However, the best evidence in Foster and Rosenzweig (2022) suggests increasing returns to land over some scale – perhaps the micropinnings for the observed indivisibilities observed.

[41] To see this, recall that θ only impacts the household's problem through the borrowing constraint, $a' \geq -\theta P k'$. Define $\hat{a} = a' + \theta P k'$ as surplus savings above the lower bound. Then combining the recursive constraints (3) and (5) to eliminate l, substituting in for $\hat{a}$, and simplifying using $r = 0$ yields the following constraints:

$$c + \hat{a} + P(1-\theta)\left(k' - k\right) = e + \hat{a} + zk^{\alpha} - \delta k$$

and

$$\hat{a} \geq 0.$$

It is clear that any change from θ to θ' is simply offset by a general equilibrium change in the price of capital from P to $P' = P(1-\theta)/(1-\theta')$, and no household's problem is affected, except that while $\hat{a}$ remains constant, the 0.41 indicates that a' declines dramatically (by 59%) to offset the change in available credit according to the equation that holds by definition, $\hat{a} = a' + \theta P k'$. The clarity of this analytical neutrality result is further justification for the $r = 0$ assumption.

poverty is perhaps as dramatic as in the credit intervention, as the fraction poor drops to just 18% of its baseline level and the probability of staying poor drops as well. These results are dramatic but not surprising: when agents face uninsured shocks, as interest rates approach the subject discount rate, savings is no longer viewed as a costly tradeoff. Consequently, they accumulate assets in unbounded fashion for the purposes of self-insurance against shocks.

When we look at the case of capital in fixed supply, although the perfect neutrality result no longer holds, the aggregate impacts are nonetheless greatly tempered. The increase in net savings is not as large in GE, but still dramatic, a 3-fold increase. As in the case of credit, the price of capital again increases, here by 61%, with the increased demand reflecting the high amount of savings. Income rises by 29%, much of it interest income. Hence, while the aggregate results are tempered, they are not completely neutral. Moreover, because it affects the accumulation of savings, the program still has important impacts on poverty. The fraction poor drops substantially even under a fixed supply of land, to 73% of its original level, and poverty is slightly less persistent with the probability of staying poor dropping by 5%.

Lastly, the combined policy shows that the combination of policies can be a positive force when capital is elastically supplied but a negative force for long run regional development. In the PE case, and comparing to the savings along scenario, we can see that credit is used (i.e., savings increases by less than in the savings only scenario). Income, consumption, and capital all increase by more than with savings alone, while poverty and the persistence of poverty are lower. When capital is elastically supplied and savings is available, access to interest-charging credit has small but positive benefits on long-run development.

However, in the case of capital in fixed supply, the small impacts on development become negative. The key to understanding the negative impact of credit is to appreciate that credit requires paying interest in this scenario. Following the above intuition for the credit only case, while the quantity of capital is completely unaffected, even the allocation of capital is largely unaffected by credit, so only the price is impacted, rising by 82% (substantially more than the 61% in the case of savings alone). But the higher price requires borrowing, and this can be seen by the lower net savings (2.17) relative to the savings alone world (2.99). This credit does not increase entrepreneurial income, since the quantity and allocation of capital is largely unchanged, but it does require interest payments which leave the regional economy, lowering income and consumption and increasing poverty and its persistence relative to the world with only savings.

In summary, the aggregate simulation results underscore that the indivisibility of capital has strong impacts on aggregate development and poverty. While one-time asset grants do not impact long run development, financial interventions can be powerful forces for development and poverty reduction in this world of indivisible investments, but this hinges critically on the elasticity of the supply of capital. General equilibrium forces can undermine positive aggregate impacts measured in partial equilibrium experiments. When capital is in fixed supply, savings is the most powerful intervention, whereas combined financial services are best when it is elastically supplied.

5.4 Aggregate Shocks?

We should also address the relevance of aggregate shocks, which we haven't modeled but our empirical analysis found to be important. We conjecture that aggregate shocks would only reinforce our main policy conclusions. Substantial aggregate shocks will increasing not just relative churning within the distribution but the level of the overall distribution itself, thereby pushing people in and out of poverty traps. This would make poverty traps, and therefore the impacts of one-time grants, even more transient. However, the impact of aggregate shocks would also hinge on whether or not the indivisible investment good is elastically supplied. In a situation where capital is in fixed supply, an aggregate shock would simply affect the price of the fixed stock (in a fashion similar to credit) with little impact on the allocation or aggregate output. In contrast, if capital is elastically supplied, it would affect the real quantity of capital, allowing productive people to either escape from or fall into temporary poverty traps, depending on the direction of the shock.

5.5 Elasticity of Supply of Land?

Given the importance of land as an indivisible investment in the empirics, a final question is whether we ought to think about semi-urban land in fixed or elastic supply. One reasonable answer is to simply assert that land captures simple area, and is in therefore in fixed supply. If the increased land value we capture reflects land purchases, so even the small-scale increases in land value among the treated come from the sales of land of the untreated, a quantitatively small but in this case economically important violation of the stable unit treatment value assumption (SUTVA) for our control. An alternative to consider is that land value, especially quality-adjusted land, is potentially elastic, however, reflecting either investments/improvements in land that increase the value and/or a true expansion of utilized land from around the perimeter of a town, for example. Following such an assumption, we can use the exogenous increase in demand of land from our experiment to calculate a back-of-the-envelope estimate of the elasticity of the supply of land. Specifically, using the results in Table 4 we estimate a percentage increase in the quantity of quality-adjusted land (calculated to be roughly 1% as a fraction of all land, participants and not). Recalling the price effects estimated in Appendix B, which implied an increase in price of 20%, we estimate a very small elasticity of $1\%/20\% = 0.05$, quite close to the inelastic benchmark, indicating the GE results are most relevant.[42]

6 Conclusions

We examine the importance of high-yield, indivisible investments in semi-urban and rural Uganda, where financial services are limited. Empirically, we have shown demand for a large-stakes-but-lower-expected-value lottery in an otherwise risk-averse population. This is especially true among those who

[42] The specifics of these calculations are also included in Appendix B.

have self-financing motives for savings and investment motives for credit. The risk-loving behavior that we observe is consistent with theories of high-yield indivisible investments, which, in principle, can lead to poverty trap dynamics and sizable gains of financial intermediation. Winners of these large lotteries invest disproportionately in indivisible assets, land in particular. We find sizable income gains for winners of the large lottery which take time to develop — they are present at 6 years but not at 4 or 18 months (though there is some evidence of consumption increases for large lottery winners at 18 months). Our theory demonstrates that given the presence of shocks, one time grants do not have lasting impacts even when poverty trap dynamics are present. When indivisible investments are inelastically supplied, the impacts of expanding financial services, especially credit, on aggregate outcomes and poverty can be substantially lower. So neither credit, grants, nor divisible land are long term policy solutions in the Ugandan semi-urban environment we study. Instead, interest-bearing savings, perhaps along the lines of Greaney, Kaboski, and Van Leemput (2016), is more important.

Substantively, the findings about land elicit questions of their own for future research and highlight the importance of work in semi-urban environments which are growing quickly and often have lagging access to infrastructure. Poverty trap dynamics have been observed for smaller investments such as livestock, which are profitable and more elastically supplied (e.g., Balboni et al. (2022)), but the importance of land merits more consideration. The development literature has focused on the interaction of finance, land titling, and investment, and also the link between land plots, overall farm size, and misallocation (e.g., Foster and Rosenzweig (2022); Gollin and Udry (2021); Acampora et al. (2022)), but the issues we raise are unique. How common are land-driven poverty traps and to what extent does land reallocation increase productivity and income? Finally, our results indicate that, even outside of major urban areas, land may be an important investment for the poor. Returns to land may result from investment in the land by more productive households or simply from holding a good that offers high capital gains, and the relative importance of each likely varies across economic environments. More research on these aspects of land markets is therefore encouraged.

References

ACAMPORA, M., L. CASABURI, AND J. WILLIS (2022): "Land rental markets: Experimental evidence from Kenya," Tech. rep., National Bureau of Economic Research.

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

BALBONI, C., O. BANDIERA, R. BURGESS, M. GHATAK, AND A. HEIL (2022): "Why do people stay poor?" *The Quarterly Journal of Economics*, 137, 785–844.

BANERJEE, A., E. BREZA, E. DUFLO, AND C. KINNAN (2019): "Can microfinance unlock a poverty trap for some entrepreneurs?" Tech. rep., National Bureau of Economic Research.

BANERJEE, A., M. FAYE, A. KRUEGER, P. NIEHAUS, AND T. SURI (2023): "Universal Basic Income: Short-Term Results from a Long-Term Experiment in Kenya," Tech. rep., UC San Diego.

BANERJEE, A. V. AND A. F. NEWMAN (1993): "Occupational choice and the process of development," *Journal of Political Economy*, 101, 274–298.

BASSI, V., R. MUOIO, T. PORZIO, R. SEN, AND E. TUGUME (2022): "Achieving scale collectively," *Econometrica*, 90, 2937–2978.

BEAMAN, L., D. KARLAN, B. THUYSBAERT, AND C. UDRY (2023): "Selection into credit markets: Evidence from agriculture in Mali," *Econometrica*, 91, 1595–1627.

BERGQUIST, L., B. FABER, T. FALLY, M. HOELZLEIN, E. MIGUEL, AND A. RODRIGUEZ-CLARE (2019): "Scaling Agricultural Policy Interventions: Theory and Evidence from Uganda," Unpublished manuscript, University of California at berkeley.

BLATTMAN, C., N. FIALA, AND S. MARTINEZ (2014): "Generating skilled self-employment in developing countries: Experimental evidence from Uganda," *The Quarterly Journal of Economics*, 129, 697–752.

BLATTMAN, C., E. P. GREEN, J. JAMISON, M. C. LEHMANN, AND J. ANNAN (2016): "The returns to microenterprise support among the ultrapoor: A field experiment in postwar Uganda," *American Economic Journal: Applied Economics*, 8, 35–64.

BRUHN, M. AND D. MCKENZIE (2009): "In Pursuit of Balance: Randomization in Practice in Development Field Experiments," *American Economic Journal: Applied Economics*, 1, 200–232.

BRYAN, G., J. DE QUIDT, M. SILVA-VARGAS, T. WILKENING, AND N. YADAV (2022): "Market Design for Land Trade: Evidence from Uganda and Kenya," *London School of Economics, London*.

BUERA, F. J. (2009): "A dynamic model of entrepreneurship with borrowing constraints: theory and evidence," *Annals of Finance*, 5, 443–464.

BUERA, F. J., J. P. KABOSKI, AND Y. SHIN (2011): "Finance and development: A tale of two sectors," *American Economic Review*, 101, 1964–2002.

——— (2014): "Macro-perspective on asset grants programs: Occupational and wealth mobility," *American Economic Review*, 104, 159–64.

——— (2020): "The Macroeconomics of Microfinance," *The Review of Economic Studies*, 88, 126–161.

BUERA, F. J. AND Y. SHIN (2013): "Financial frictions and the persistence of history: A quantitative exploration," *Journal of Political Economy*, 121, 221–272.

CARTER, M. R. AND T. J. LYBBERT (2012): "Consumption versus asset smoothing: testing the implications of poverty trap theory in Burkina Faso," *Journal of Development Economics*, 99, 255–264.

CASABURI, L. AND R. MACCHIAVELLO (2019): "Demand and supply of infrequent payments as a commitment device: evidence from Kenya," *American Economic Review*, 109, 523–555.

CAUNEDO, J. AND N. KALA (2021): "Mechanizing Agriculture," Working Paper 29061, National Bureau of Economic Research.

CHERNOZHUKOV, V., I. FERNÁNDEZ-VAL, AND B. MELLY (2013): "Inference on counterfactual distributions," *Econometrica*, 81, 2205–2268.

CLONTS, H. A. (1970): "Influence of urbanization on land values at the urban periphery," *Land Economics*, 46, 489–497.

CROSON, R. AND U. GNEEZY (2009): "Gender Differences in Preferences," *Journal of Economic Literature*, 47, 448–74.

DE MEL, S., D. MCKENZIE, AND C. WOODRUFF (2008): "Returns to capital in microenterprises: evidence from a field experiment," *The Quarterly Journal of Economics*, 123, 1329–1372.

——— (2014): "Business training and female enterprise start-up, growth, and dynamics: Experimental evidence from Sri Lanka," *Journal of Development Economics*, 106, 199–210.

DEININGER, K. W. ET AL. (2003): *Land policies for growth and poverty reduction*, vol. 41181, World Bank Publications.

DONOVAN, K. (2021): "The equilibrium impact of agricultural risk on intermediate inputs and aggregate productivity," *The Review of Economic Studies*, 88, 2275–2307.

EGGER, D., J. HAUSHOFER, E. MIGUEL, P. NIEHAUS, AND M. W. WALKER (2019): "General equilibrium effects of cash transfers: experimental evidence from Kenya," Tech. rep., National Bureau of Economic Research.

FOSTER, A. D. AND M. R. ROSENZWEIG (2022): "Are there too many farms in the world? labor market transaction costs, machine capacities, and optimal farm size," *Journal of Political Economy*, 130, 636–680.

GALOR, O. AND J. ZEIRA (1993): "Income distribution and macroeconomics," *The Review of Economic Studies*, 60, 35–52.

GERTLER, P., S. HIGGINS, A. SCOTT, AND E. SEIRA (2023): "Using Lotteries to Attract Deposits," Tech. rep., National Bureau of Economic Research.

GOCHBERG, W. (2021): "The social costs of titling land: Evidence from Uganda," *World Development*, 142, 105376.

GOLLIN, D. AND C. UDRY (2021): "Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture," *Journal of Political Economy*, 129, 1–80.

GOODMAN-BACON, A. (2021): "The long-run effects of childhood insurance coverage: Medicaid implementation, adult health, and labor market outcomes," *American Economic Review*, 111, 2550–2593.

GREANEY, B. P., J. P. KABOSKI, AND E. VAN LEEMPUT (2016): "Can self-help groups really be "self-help"?" *The Review of Economic Studies*, 83, 1614–1644.

HAUSHOFER, J. AND J. SHAPIRO (2016): "The short-term impact of unconditional cash transfers to the poor: experimental evidence from Kenya," *The Quarterly Journal of Economics*, 131, 1973–2042.

——— (2018): "The long-term impact of unconditional cash transfers: experimental evidence from Kenya," Tech. rep., Busara Center for Behavioral Economics, Nairobi, Kenya.

HERSKOWITZ, S. (2021): "Gambling, saving, and lumpy liquidity needs," *American Economic Journal: Applied Economics*, 13, 72–104.

HOLT, C. A. AND S. K. LAURY (2002): "Risk aversion and incentive effects," *American Economic Review*, 92, 1644–1655.

HSIEH, C.-T. AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 124, 1403–1448.

KABOSKI, J. P. AND R. M. TOWNSEND (2011): "A structural evaluation of a large-scale quasi-experimental microfinance initiative," *Econometrica*, 79, 1357–1406.

KAPLAN, G. AND G. L. VIOLANTE (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82, 1199–1239.

KAPLAN, G., G. L. VIOLANTE, AND J. WEIDNER (2014): "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 45, 77–153.

KORAH, P. I., T. MATTHEWS, AND D. TOMERINI (2019): "Characterising spatial and temporal patterns of urban evolution in Sub-Saharan Africa: The case of Accra, Ghana," *Land Use Policy*, 87, 104049.

LAFAVE, D. AND D. THOMAS (2016): "Farms, Families, and Markets: New Evidence on Completeness of Markets in Agricultural Settings," *Econometrica*, 84, 1917–1960.

LAGAKOS, D., M. MOBARAK, AND M. WAUGH (2018): "The Welfare Effects of Encouraging Rural-Urban Migration," Working Papers 2018–002, Human Capital and Economic Opportunity Working Group.

LEE, D. S. (2009): "Training, wages, and sample selection: Estimating sharp bounds on treatment effects," *The Review of Economic Studies*, 76, 1071–1102.

LYBBERT, T. J. AND C. B. BARRETT (2011): "Risk-Taking Behavior in the Presence of Nonconvex Asset Dynamics," *Economic Inquiry*, 49, 982–988.

LYBBERT, T. J., C. B. BARRETT, S. DESTA, AND D. L. COPPOCK (2004): "Stochastic Wealth Dynamics and Risk Management among a Poor Population," *The Economic Journal*, 114, 750–777.

MCKENZIE, D. J. AND C. WOODRUFF (2006): "Do Entry Costs Provide an Empirical Basis for Poverty Traps? Evidence from Mexican Microenterprises," *Economic Development and Cultural Change*, 55, 3–42.

MEISSNER, T., X. GASSMANN, C. FAURE, AND J. SCHLEICH (2023): "Individual characteristics associated with risk and time preferences: A multi country representative survey," *Journal of Risk and Uncertainty*, 66, 77–107.

MIDRIGAN, V. AND D. Y. XU (2014): "Finance and misallocation: Evidence from plant-level data," *American Economic Review*, 104, 422–58.

PIKETTY, T. (1997): "The dynamics of the wealth distribution and the interest rate with credit rationing," *The Review of Economic Studies*, 64, 173–189.

RESTUCCIA, D. AND R. ROGERSON (2008): "Policy distortions and aggregate productivity with heterogeneous establishments," *Review of Economic Dynamics*, 11, 707–720.

ROSENZWEIG, M. R. AND H. P. BINSWANGER (1993): "Wealth, weather risk and the composition and profitability of agricultural investments," *The economic journal*, 103, 56–78.

ROSENZWEIG, M. R. AND C. UDRY (2020): "External Validity in a Stochastic World: Evidence from Low-Income Countries," *The Review of Economic Studies*, 87, 343–381.

ROTHSCHILD, M. AND J. STIGLITZ (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information*," *The Quarterly Journal of Economics*, 90, 629–649.

TIBSHIRANI, R. (1996): "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society: Series B (Methodological)*, 58, 267–288.

WINEMAN, A. AND T. S. JAYNE (2018): "Land prices heading skyward? An analysis of farmland values across Tanzania," *Applied Economic Perspectives and Policy*, 40, 187–214.

Figures

Figure 1: **Wealth-Dependent Behavior and Poverty Traps**

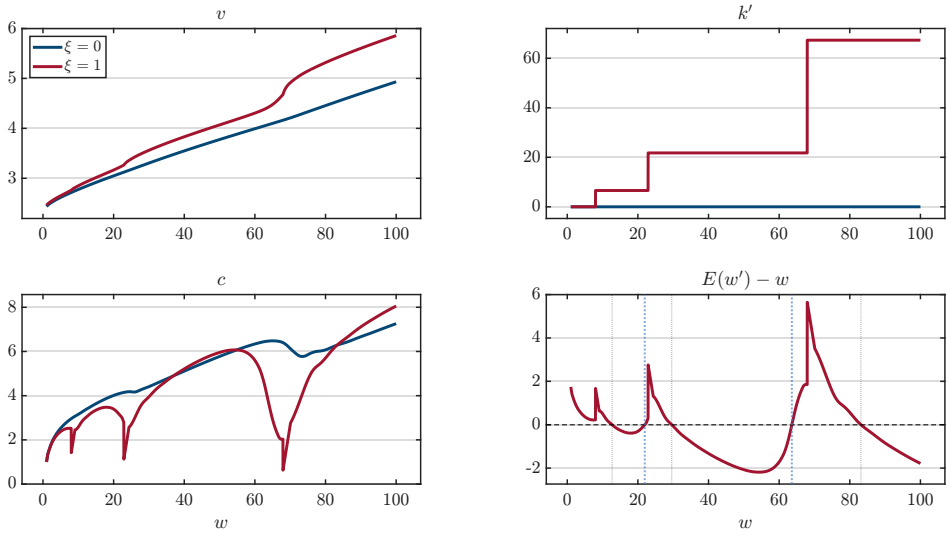

Notes: The four panels plot policy functions, the value function, and expected wealth dynamics from the calibrated model as functions of beginning-of-period wealth w. The upper-left panel shows the value function $v(w)$ in consumption-equivalent units; the upper-right panel shows the capital choice, $k'(w)$; the bottom-left panel shows consumption $c(w)$; and the bottom-right panel shows the expected change in wealth $E(w') - w$. The blue and red lines correspond to the $\xi = 0$ and the $\xi = 1$ states, respectively; the bottom-right panel is plotted only for $\xi = 1$. In that panel, the blue dotted vertical lines mark points for which $E(w') - w$ reaches zero from below, and the black dotted vertical lines mark points for which $E(w') - w$ reaches zero from above.

Figure 2: **Lottery Choices and Indivisibilities**

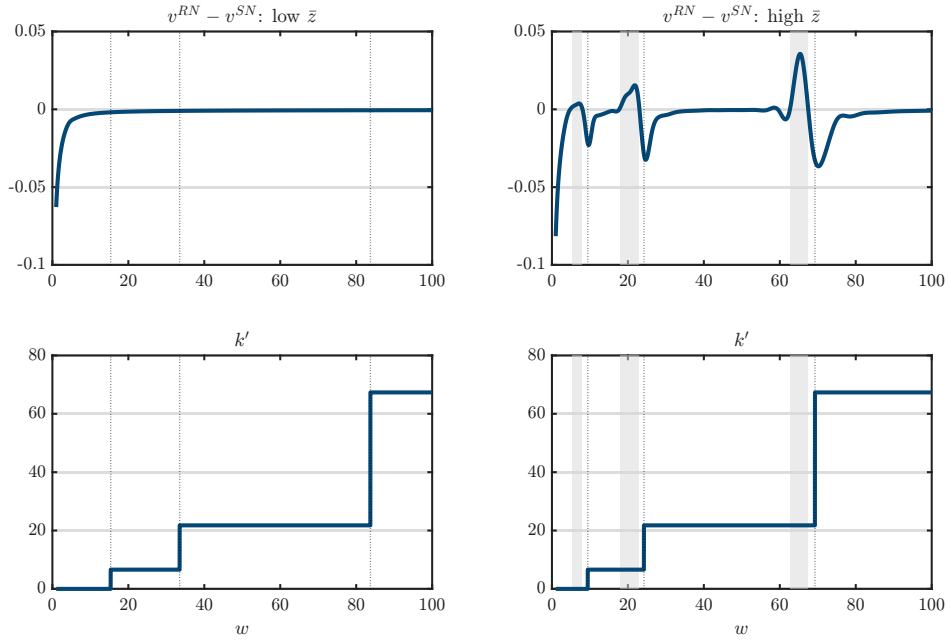

Notes: The four panels plot lottery preferences and capital choices in the calibrated model as functions of beginning-of-period wealth w, with the left column corresponding to households with low permanent entrepreneurial productivity $\bar{z}$ and the right column to high $\bar{z}$. The top panels show the difference in value between the risky and safe lotteries, $v^{RN}(w) - v^{SN}(w)$, and the bottom panels show the corresponding next-period capital choice $k'(w)$. Gray-shaded regions indicate wealth levels at which the risky lottery is preferred ($v^{RN} > v^{SN}$). Dotted vertical lines mark discrete jumps in $k'(w)$.

Figure 3: **Experimental Timeline, Experimental Design, and Lottery Choices**

(a) Timeline

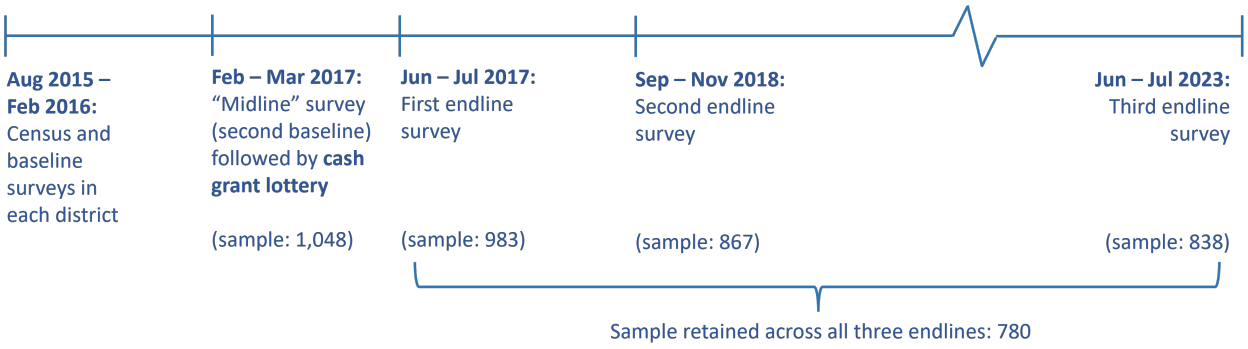

(b) Experimental Design and Lottery Choices

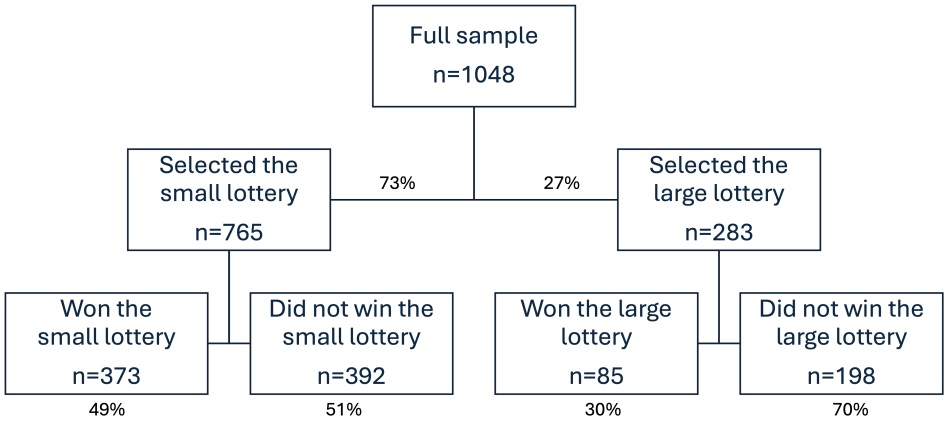

Figure 4: **Distribution Regressions: Effect of Winning the Lottery on the Distributions of Income at Each Endline**

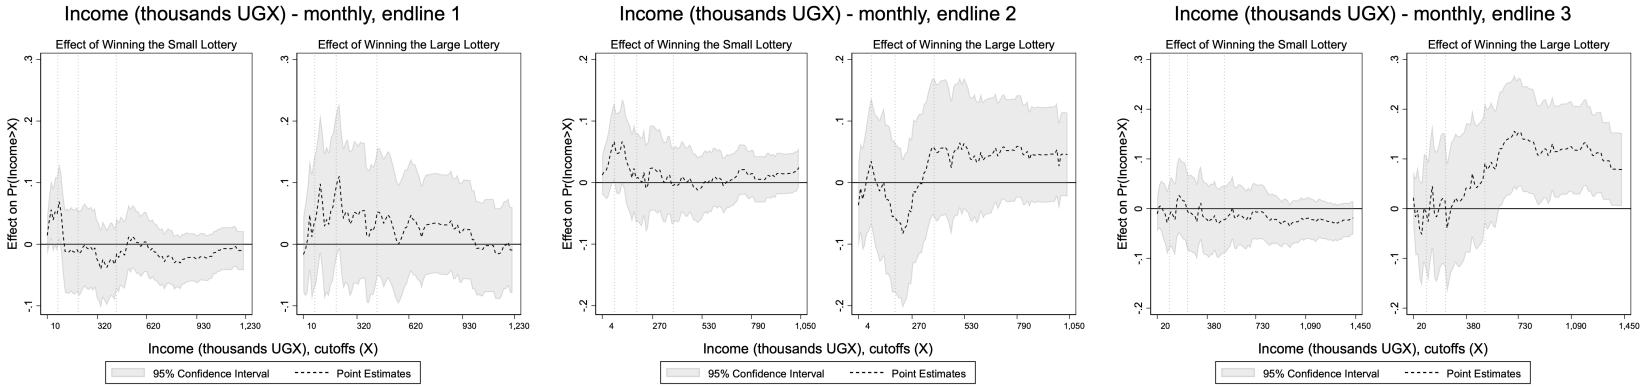

Notes: Vertical dotted lines correspond to the 25th, 50th, and 75th percentiles of the outcome, respectively.

Figure 5: **Simulated Long-Run Treatment Effects on Income and Capital**

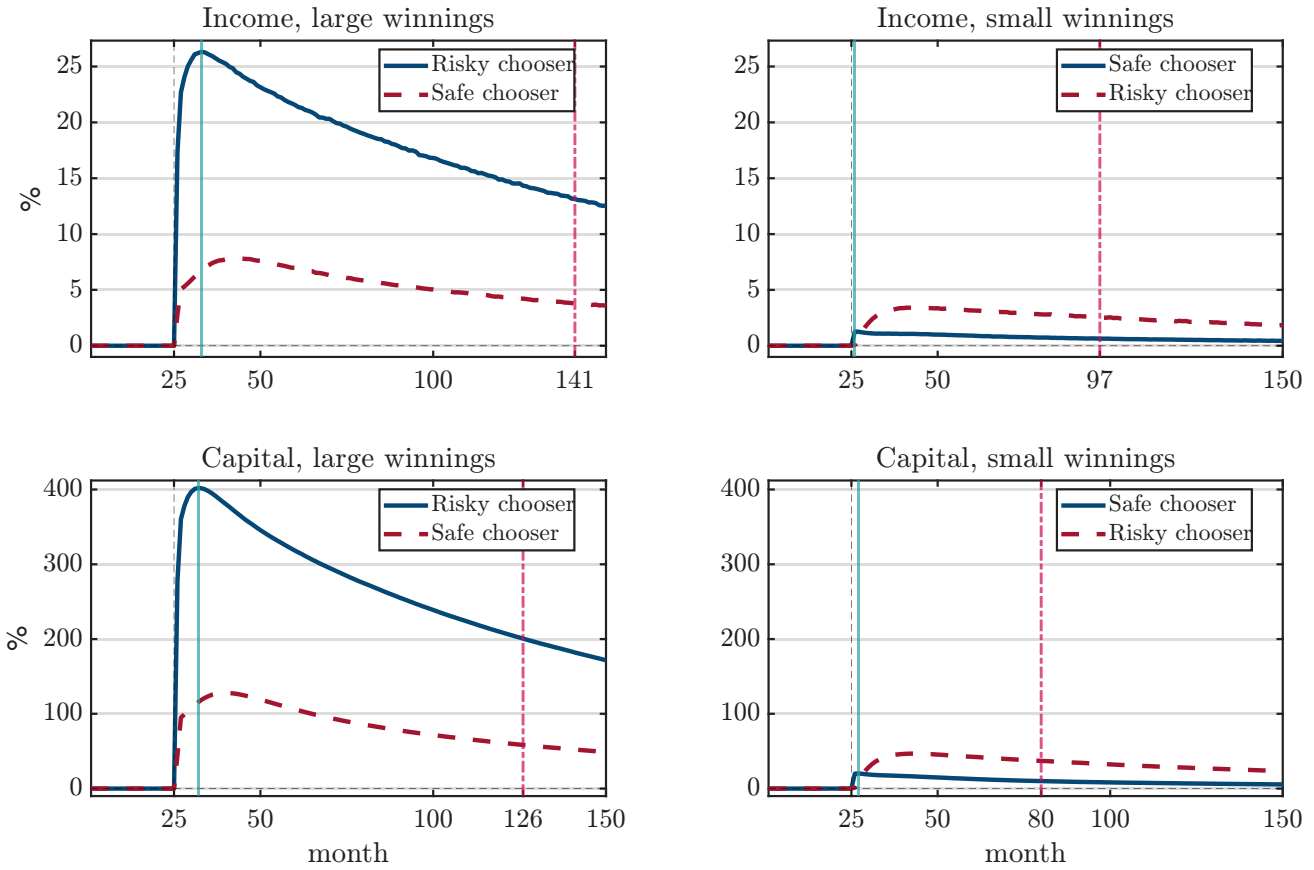

Notes: These plot the simulated treatment effects in the calibrated model, i.e., for both treatment and control, we calculate the average change from midline in the averages as a percentage of the average midline level, and then difference these across treatment (lottery winners) and control (lottery losers). The solid blue lines are for the actual (preferred) lottery chosen, while the dashed red lines are hypothetical treatment effects if people had been granted the opposite (non-preferred) lottery. The vertical dashed black line indicates the time of treatment (midline). The vertical solid blue line indicates the timing of the peak treatment effect, while the vertical dashed pink line indicates the time at which treatment has fallen to half its peak. The time between the two is the half-life of the impact.

Tables

Table 1: **Balance between small lottery winners and non-winners**

	won small lottery	N	did not win small lottery	N	diff	p-value
Household budget components						
monthly household income$_m$	361,708	293	389,628	276	-27,919	0.369
weekly consumption$_m$	40,291	293	39,176	276	1,114	0.629
savings$_m$	287,993	293	263,620	276	24,374	0.465
credit outstanding$_m$	193,427	293	212,588	276	-19,161	0.575
home durable value$_b$	537,797	293	578,025	276	-40,229	0.586
Investment categories						
total divisible investments$_m$	404,246	293	342,832	276	61,415	0.223
small livestock and ag. assets$_b$	190,005	293	177,030	276	12,975	0.536
bus inventory$_m$	167,311	293	148,520	276	18,791	0.557
total indivisible investments$_m$	11,414,571	293	11,652,123	276	-237,551	0.849
large livestock and ag. assets$_b$	346,359	293	259,962	276	86,397	0.281
land value$_b$	11,024,083	293	11,510,133	276	-486,049	0.697
bus assets, no stock$_m$	421,364	293	282,469	276	138,894*	0.072
Other financial indicators						
operates non-farm business$_m$ (0/1)	.54	293	.55	276	-.0079	0.851
farmer$_m$ (0/1)	.75	293	.77	276	-.017	0.625
work hours per week$_m$	77	293	77	276	-.11	0.955
had negative shock since baseline$_m$ (0/1)	.63	293	.69	276	-.057	0.152
has formal savings$_m$ (0/1)	.11	293	.15	276	-.043	0.128
acquired loans since baseline$_m$ (0/1)	.33	293	.34	276	-.013	0.744
Demographic characteristics						
female (0/1)	.48	293	.52	276	-.041	0.335
household head (0/1)	.6	293	.6	276	.0026	0.949
respondent age	36	293	37	276	-1.5	0.103
education beyond primary school (0/1)	.26	293	.29	276	-.03	0.416
num people in household$_b$	5.1	293	5	276	.086	0.647
num adult females$_b$	1.1	293	1.2	276	-.082*	0.079
num adult males$_b$	1.5	293	1.4	276	.048	0.591
num children$_b$	2.5	293	2.4	276	.12	0.388
Observations	569					

Notes: All quantities in UGX; Outliers top/bottom coded to 95th/5th percentile; Midline denoted by **m** and baseline denoted by **b**. Sample consists of those retained in all three endline surveys (those used in our analysis of treatment effects). **bus** is an abbreviation for business; * $p < 0.1$, ** $p < .05$, *** $p < 0.01$.

Table 2: **Balance between large lottery winners and non-winners**

	won large lottery	N	did not win large lottery	N	diff	p-value
Household budget components						
monthly household income$_m$	431,809	68	373,017	143	58,793	0.299
weekly consumption$_m$	45,929	68	43,518	143	2,410	0.587
savings$_m$	364,206	68	293,531	143	70,674	0.232
credit outstanding$_m$	178,529	68	207,126	143	-28,596	0.642
home durable value$_b$	756,398	68	775,288	143	-18,889	0.902
Investment categories						
total divisible investments$_m$	618,654	68	454,441	143	164,214	0.128
small livestock and ag. assets$_b$	196,756	68	221,050	143	-24,294	0.560
bus inventory$_m$	301,691	68	197,399	143	104,293	0.133
total indivisible investments$_m$	18,558,235	68	15,121,692	143	3,436,543	0.186
large livestock and ag. assets$_b$	638,154	68	650,272	143	-12,118	0.954
land value$_b$	17,291,038	68	14,146,459	143	3,144,579	0.219
bus assets, no stock$_m$	542,794	68	513,441	143	29,354	0.860
Other financial indicators						
operates non-farm business$_m$ (0/1)	.62	68	.62	143	.0023	0.975
farmer$_m$ (0/1)	.75	68	.73	143	.023	0.728
work hours per week$_m$	77	68	77	143	-.064	0.986
had negative shock since baseline$_m$ (0/1)	.6	68	.62	143	-.012	0.863
has formal savings$_m$ (0/1)	.059	68	.091	143	-.032	0.426
acquired loans since baseline$_m$ (0/1)	.24	68	.34	143	-.1	0.140
Demographic characteristics						
female (0/1)	.37	68	.43	143	-.059	0.418
household head (0/1)	.74	68	.64	143	.099	0.155
respondent age	39	68	38	143	.91	0.596
education beyond primary school (0/1)	.31	68	.2	143	.11*	0.091
num people in household$_b$	5.4	68	5.4	143	-.015	0.967
num adult females$_b$	1.2	68	1.1	143	.057	0.549
num adult males$_b$	1.6	68	1.4	143	.21	0.199
num children$_b$	2.6	68	2.9	143	-.28	0.289
Observations	211					

Notes: All quantities in UGX; Outliers top/bottom coded to 95th/5th percentile; Midline denoted by **m** and baseline denoted by **b**. Sample consists of those retained in all three endline surveys. **bus** is an abbreviation for business; * $p < 0.1$, ** $p < .05$, *** $p < 0.01$.

Table 3: **Statistically significant differences between those choosing the large v. small lottery**

	chose large mean	N	chose small mean	N	Diff	p-value
Income and Consumption						
monthly crop income$_m$*	100,208	283	71,148	765	29,061***	0.000
Δ ln monthly household income$_{m\text{-}b}$*	-.34	283	.14	765	-.48	0.023
crop income/household income$_m$	.3	283	.25	765	.05	0.009
monthly crop income/adult equiv$_m$*	30,562	283	22,981	765	7,580	0.001
ln monthly crop income$_m$*	12	283	12	765	.21	0.000
Δ ln monthly bus income$_{m\text{-}b}$*	6.2	283	7.2	765	-.96	0.024
weekly consumption$_m$	44,466	283	39,453	765	5,013	0.011
ln weekly consumption$_m$	10	283	10	765	.18	0.013
Δ ln weekly consumption$_{m\text{-}b}$*	-.55	283	-.69	765	.14	0.039
Savings and Wealth						
savings$_m$	322,675	283	275,817	765	46,858	0.095
Δ ln savings$_{m\text{-}b}$	1.9	283	.94	765	.91	0.029
bus assets$_m$	824,954	283	577,814	765	247,140	0.011
bus assets/wealth$_m$	.29	283	.24	765	.053	0.047
bus assets/adult equiv$_m$	295,958	283	218,673	765	77,285	0.069
ln bus assets$_m$	5.3	283	4.3	765	.99	0.031
wealth (sav + bus assets)$_b$	1,245,155	283	920,185	765	324,970	0.006
wealth/adult equiv$_m$	431,561	283	337,778	765	93,784	0.066
ln wealth$_m$	11	283	11	765	.65	0.048
Δ ln wealth$_{m\text{-}b}$*	1.3	283	.61	765	.7	0.061
net wealth (sav + bus assets - credit)$_m$	1,070,910	283	744,418	765	326,492	0.006
net wealth/adult equiv$_m$	377,813	283	277,547	765	100,265	0.049
ln net wealth$_m$	16	283	16	765	.075	0.060
land value$_b$	14,579,702	283	10,764,138	765	3,815,564	0.000
land value/adult equiv$_b$	4,300,681	283	3,288,978	765	1,011,702	0.003
ln land value$_b$	15	283	14	765	.99	0.045
Desire to Invest						
wants credit to increase income$_b$* (0/1)	.84	283	.78	765	.062	0.027
would invest >\$100$_b$* (0/1)	.95	283	.91	765	.038	0.040
would use credit for bus investment$_b$* (0/1)	.67	283	.59	765	.073	0.031
Demographic Characteristics						
female (0/1)*	.42	283	.51	765	-.09	0.009
household head (0/1)	.66	283	.6	765	.058	0.085
respondent age	37	283	35	765	2.3	0.003
num people in household$_b$	5.5	283	5	765	.47	0.004
num adult males$_b$	1.5	283	1.4	765	.14	0.076
num children$_b$*	2.8	283	2.5	765	.34	0.004
Observations	1,048					

*Denotes that variable was also selected by lasso. All quantities in UGX. Outliers top/bottom coded to 95th/5th percentile. Midline denoted by **m** and baseline denoted by **b**; **bus** is an abbreviation for business. Full list of covariates on which we test for differences between those selecting the small versus large lottery is in Table D.1. Complete list of lasso-selected covariates is in Table D.2.

Table 4: **Grant effects on investment**

		Productive divisible investments			Productive indivisible investments			
	Home durables	Tot. div.	Small livestock and ag. tools	Bus. inventory	Tot. indiv.	Large livestock and ploughs	Bus. durables	Land
Endline 1: 4 Months								
won lottery (0/1)	110,302	240,566***	86,434***	155,398***	-468,154	124,591**	85,337	-669,112
	(67,086)	(64,722)	(29,422)	(43,463)	(1,015,458)	(52,825)	(71,990)	(1,058,630)
won large lottery (0/1)	-123,607	15,509	-97,716	156,821	4,426,545**	108,340	23,599	4,856,940**
	(133,220)	(149,842)	(63,938)	(108,106)	(2,183,416)	(125,458)	(179,480)	(2,252,993)
$\beta_1 + \beta_2$	-13,304	256,075	-11,282	312,220	3,958,391	232,931	108,937	4,187,828
P-value: $\beta_1 + \beta_2 = 0$	0.91	0.06	0.84	0.00	0.04	0.04	0.50	0.04
Endline 2: 18 Months								
won lottery (0/1)	32,487	90,743	35,965	70,193*	1,156,837	109,986**	131,873	809,046
	(70,087)	(60,299)	(27,389)	(42,347)	(1,501,866)	(52,901)	(117,593)	(1,512,409)
won large lottery (0/1)	-64,581	-81,262	-45,630	-6,253	6,825,268*	-126,973	233,171	7,167,692**
	(148,575)	(135,123)	(61,103)	(101,313)	(3,541,022)	(116,586)	(294,098)	(3,472,755)
$\beta_1 + \beta_2$	-32,094	9,481	-9,665	63,941	7,982,105	-16,987	365,044	7,976,738
P-value: $\beta_1 + \beta_2 = 0$	0.81	0.94	0.86	0.49	0.01	0.87	0.17	0.01
Endline 3: 6 years								
won lottery (0/1)	25,621	11,582	33,939	4,592	1,198,692	-18,196	41,449	1,479,878
	(91,417)	(96,520)	(35,785)	(72,545)	(2,590,381)	(55,856)	(140,791)	(2,508,803)
won large lottery (0/1)	-50,733	75,169	-3,959	64,588	-855,056	188,292	60,284	-1,029,594
	(181,015)	(211,973)	(78,835)	(155,997)	(5,531,055)	(120,695)	(322,087)	(5,288,633)
$\beta_1 + \beta_2$	-25,112	86,752	29,980	69,180	343,636	170,096	101,733	450,284
P-value: $\beta_1 + \beta_2 = 0$	0.87	0.65	0.67	0.62	0.94	0.11	0.73	0.92
Observations	780	780	780	780	780	780	780	780
Control mean if risk loving = 0	909,277	835,857	362,871	413,703	33,261,825	234,304	779,687	31,418,790
Control mean if risk loving = 1	977,051	954,174	441,906	427,626	40,124,778	265,566	958,975	38,069,914

Notes: All quantities in UGX; Outliers top/bottom coded to 95th/5th percentile; Heteroskedasticity-robust standard errors in parentheses. Variables are in constant 2017 UGX. Controls include pre-treatment levels of outcome, hh income, patience, gender, hh head, age, age^2, num ad females, num ad males, num children, district FE's. ***1%, **5%, *10%.

Table 5: **Effect of winning the grant on land purchases, sales, and investment at third endline**

	(1) purchased land between midline & endline3	(2) invested in land between midline & endline3	(3) sold land between midline & endline3	(4) added land to an existing plot between 2016 & endline3	(5) added land to least cultivated plot between 2016 & endline3	(6) added land to most cultivated plot between 2016 & endline3	(7) number of plots at endline3	(8) average plot size at endline3
won lottery (0/1)	.011	.032	-.021	-.03	-.021**	-.016	.07	.022
	(.038)	(.039)	(.031)	(.027)	(.0098)	(.022)	(.13)	(.098)
won large lottery (0/1)	.15**	-.026	-.04	.092	.049**	.0014	-.2	-.14
	(.078)	(.081)	(.065)	(.059)	(.022)	(.046)	(.25)	(.18)
$\beta_1 + \beta_2$	.16	.0057	-.061	.062	.028	-.014	-.13	-.12
P-value: $\beta_1 + \beta_2 = 0$	.017	.94	.29	.25	.17	.72	.51	.45
Control mean if risk loving = 0	.37	.36	.18	.15	.023	.093	2.2	1.1
Control mean if risk loving = 1	.33	.41	.22	.13	0	.09	2.6	1.2
R^2	.08	.043	.031	.047	.02	.039	.095	.15
Observations	838	838	838	838	838	838	838	838

Notes: Standard errors in parentheses. Outcomes for columns 1-3 are dummy variables for between the midline and the third endline. Column 4's dummy variable includes 2016 to the third endline. Columns 5-6 are measured at the third endline. Columns 6-7 are dummies for adding land between 2016 and the third endline to the least or most cultivated plot at the third endline. Controls include: pre-treatment income, patience, gender, hh head status, age, age^2, num ad females, num ad males, num children. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: **Grant effects on other components of the household budget constraint**

	(1) monthly consumption	(2) monthly household income	(3) savings	(4) net credit
Endline 1: 4 Months				
won lottery (0/1)	4,910	4,587	71,718**	-3,018
	(12,387)	(21,545)	(32,236)	(19,435)
won large lottery (0/1)	19,141	22,085	88,226	-18,248
	(24,058)	(47,615)	(73,404)	(42,759)
$\beta_1 + \beta_2$	24,051	26,671	159,944	-21,266
P-value: $\beta_1 + \beta_2 = 0$	0.25	0.53	0.02	0.58
Endline 2: 18 Months				
won lottery (0/1)	18,435	9,796	49,446	-28,923
	(12,414)	(18,951)	(32,992)	(24,630)
won large lottery (0/1)	25,540	1,566	48,237	19,065
	(24,662)	(41,747)	(69,110)	(52,839)
$\beta_1 + \beta_2$	43,975	11,361	97,683	-9,857
P-value: $\beta_1 + \beta_2 = 0$	0.04	0.76	0.11	0.83
Endline 3: 6 years				
won lottery (0/1)	584	-1,061	28,424	-45,036
	(16,033)	(23,047)	(44,808)	(38,357)
won large lottery (0/1)	39,378	90,760*	2,341	132,215
	(33,183)	(50,360)	(110,806)	(80,593)
$\beta_1 + \beta_2$	39,962	89,699	30,765	87,178
P-value: $\beta_1 + \beta_2 = 0$	0.17	0.05	0.76	0.22
Observations	780	780	780	780
Control mean if risk loving = 0	176,293	389,628	263,620	212,588
Control mean if risk loving = 1	195,832	373,017	293,531	207,126

Notes: Standard errors in parentheses. All quantities in UGX. Outliers top/bottom coded to 95th/5th percentile. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: **Decomposing the impact on income: Fraction of the effect on total income accounted for by each sub-source of income**

	Business Income	Crop Income	Livestock Income	Salary and Wage Income	Remittance Income
Large lottery winners	0.27	0.49**	0.01	0.12	0.11
	(0.22)	(0.23)	(0.04)	(0.11)	(0.11)

Notes: In this table, we estimate the extent to which changes in each subcategory of income account for the effect of winning the large lottery on total household income; i.e., the interpretation of the first element of the table is that a change (in this case, an increase) in business income accounts for 27% of the effect of winning the large lottery on total household income. The columns mechanically sum to 1.00, as this is the all-inclusive set of components of household income. Each point estimate is a non-linear combination of coefficients across equations. Specifically, we divide the effect of winning the lottery on the non-winsorized sub-component of income (e.g., business income) by the effect of winning the lottery on non-winsorized *total* household income (with each estimated in separate equations; we do not winsorize the outcomes, as we do in our other analyses, because winsorizing each outcome separately means that they will not sum up). Standard errors are in parentheses and are estimated using the delta method. ***1%, **5%, *10%.

Table 8: **Crop income response to land acquisitions**

	Monthly Crop Income$_{e3}$ (UGX)	
	(1)	(2)
added land for cultivation to an existing plot$_{e3}$	26,088**	
	(11,926)	
added land to least cultivated plot$_{e3}$		27,494
		(34,566)
number of plots for cultivation$_{e3}$	-96,008*	-93,237*
	(53,567)	(54,328)
acres cultivated$_{e3}$	25,132***	26,005***
	(2,002)	(1,964)
Observations	757	757
R^2	0.288	0.284
Outcome Mean	99,924	99,924
Outcome SD	126,111	126,111

Notes: Standard errors in parentheses. Sample is restricted to individuals with a non-zero amount of plots. Controls include: whether the respondent farmed at midline, level of the outcome at midline, household income at midline, patience, gender, hh head status, age, age^2, num ad females, num ad males, num children. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: **Calibration and targeted moments**

	A. Assigned Parameters	
r	interest rate	0.00
θ	borrowing limit	0.00
δ	depreciation rate	0.01
p	survival probability	$1 - 1/480$
	B. Calibrated Parameters	
β	discount factor	0.966
σ_e	s.d. ε shock	0.795
σ_z	s.d. v shock	1.210
$k_{\min}$	minimum scale of indivisible capital	6.593
$k_{\max}$	maximum scale of indivisible capital	143.25
$\bar{z}$	mean entrepreneurial productivity	0.036
ν_z	dispersion of permanent entrepreneurial productivity	0.254
ν_e	dispersion of permanent labor income	0.702
p_L	persistence of non-entrepreneurial state	0.953
p_H	persistence of entrepreneurial state	0.986

C. Moments

C.1 Distributions and Lottery Choices

	Model	Data		Model	Data
Income			*Capital Values*		
P10	0.218	0.234	P10	0.000	0.000
P25	0.450	0.442	P25	0.000	3.033
P50	1.000	1.000	P50	4.943	15.163
P75	2.232	1.942	P75	16.325	43.972
P90	4.882	3.457	P90	50.474	90.976
Mean	2.355	1.440	Mean	12.693	32.752
Liquid Assets			*Lottery Choices*		
P10	0.139	0.058	Choose Risky	0.27	0.27
P25	0.571	0.378	Choose Safe	0.73	0.73
P50	1.782	1.466	$\sigma^{\S}$	0.007	—
P75	5.042	5.942	Choose Risky (no error)	0.207	—
P90	16.230	15.278	Frac. Risky due to error	0.285	—
Mean	7.207	4.611			

Notes: Distributions are normalized by median income. $^{\S}\sigma$ is the scale parameter of the logit error in lottery choice. "Choose Risky (no error)" is the fraction choosing the risky lottery when $\sigma = 0$. "Frac. Risky due to error" is the share of risky choosers whose choice is driven by the logit error rather than the value difference. Table continues on next page.

Table 9: **Calibration and targeted moments** (*continued*)

C. Moments (*continued*)

C.2 Lottery Treatment Effects

	Choose Safe Lottery				Choose Risky Lottery			
	Non-Winner		Winner		Non-Winner		Winner	
	Model	Data	Model	Data	Model	Data	Model	Data
Population Share	0.374	0.374	0.356	0.356	0.189	0.189	0.081	0.081
Baseline Shares								
Income	0.365	0.375	0.348	0.325	0.201	0.205	0.086	0.095
Capital	0.352	0.350	0.335	0.317	0.220	0.224	0.094	0.110
Liquid Assets°	0.356	0.299	0.339	0.331	0.213	0.259	0.091	0.111
Change in Shares (Endline 1)°								
Income	0.001	−0.009	0.002	0.004	−0.008	−0.005	0.005	0.011
Capital	−0.009	0.009	−0.004	−0.007	−0.008	−0.017	0.020	0.015
Liquid Assets	−0.029	0.023	−0.012	0.021	0.021	−0.056	0.021	0.012
Change in Shares (Endline 2)								
Income	−0.003	−0.024	−0.002	0.012	−0.002	−0.013	0.007	0.025
Capital	−0.011	−0.003	−0.007	0.015	−0.002	−0.028	0.020	0.016
Liquid Assets°	−0.014	−0.003	−0.009	0.040	0.007	−0.051	0.016	0.013
Change in Shares (Endline 3)°								
Income	−0.000	−0.014	0.001	0.008	−0.004	−0.005	0.004	0.011
Capital	−0.004	−0.001	−0.002	0.010	−0.005	0.001	0.011	−0.010
Liquid Assets	−0.009	0.047	−0.007	0.011	0.005	−0.050	0.011	−0.008

C.3 Tercile Transition Matrices†

Income

Model

		Current		
		Low	Med	High
Future	Low	0.569	0.303	0.128
	Medium	0.303	0.400	0.297
	High	0.128	0.296	0.575

Data

		Current		
		Low	Med	High
Future	Low	0.681	0.214	0.104
	Medium	0.236	0.588	0.176
	High	0.082	0.198	0.720

Capital

Model

		Current		
		Low	Med	High
Future	Low	0.640	0.220	0.140
	Medium	0.245	0.549	0.206
	High	0.115	0.231	0.654

Data

		Current		
		Low	Med	High
Future	Low	0.679	0.281	0.111
	Medium	0.229	0.438	0.230
	High	0.093	0.281	0.659

Notes: °Not targeted in calibration; reported for comparison. All other moments in C.1, C.2, and C.3 are targeted. †In C.3, future income tercile is 4 months after the current tercile, and future capital tercile is 58 months after the current tercile.

Table 10: **Counterfactual aggregate impacts**

	Benchmark	Divisible Capital		Credit $\theta=0.25$, $r=0$		Savings $\theta=0$, $r=0.03$		Combined $\theta=0.25$, $r=0.03$	
		PE	GE	PE	GE	PE	GE	PE	GE
Agg. Consumption	1.00	1.82	1.01	1.56	1.00	2.77	1.29	2.86	1.23
Agg. Income	1.00	1.84	1.01	1.56	1.00	2.82	1.31	2.91	1.24
Agg. Net Savings	1.00	1.66	0.59	0.89	0.41	8.28	2.99	7.19	2.17
Agg. Capital	1.00	3.35	*1.00*	2.51	*1.00*	3.96	*1.00*	4.58	*1.00*
Value of Capital	1.00	3.35	1.59	2.51	1.33	3.96	1.61	4.58	1.82
Price of Capital	1.00	*1.00*	1.59	*1.00*	1.33	*1.00*	1.61	*1.00*	1.82
Capital Income	1.00	3.00	1.08	2.27	1.00	3.32	1.06	3.76	1.09
Fraction Poor	1.00	0.38	1.17	0.47	1.00	0.18	0.73	0.13	0.83
Prob. Stay Poor	1.00	0.92	0.96	0.96	1.00	0.90	0.95	0.86	0.93

Notes: Values are reported relative to the benchmark, which has been normalized to one. Initial benchmark values are aggregate consumption = 3.09, aggregate income = 3.14, aggregate net savings = 9.61, aggregate capital = 16.93, value of capital = 16.93, price of capital = 1.00, capital income = 1.44, fraction poor = 0.33, and probability of staying poor = 0.90. Households whose wealth falls in the bottom 33.33% of the benchmark wealth distribution are classified as poor, and the probability of staying poor is calculated 4 months later, conditional on being poor initially. Divisible Capital allows households to purchase any value of $k \in [0, k_{\max}]$. PE fixes the price of capital to 1. GE fixes the aggregate quantity of capital to that in the benchmark. In both cases, values that are constant by construction are indicated by italics.