The Macroeconomics of Large-Scale Cash Transfers

This is a preliminary draft. Please click here for the most recent version.

Francisco J. Buera[a], Yongseok Shin[b], Kuldeep Singh[c]

[a] *Washington University in St. Louis and NBER, St. Louis, USA*
[b] *Washington University in St. Louis, Federal Reserve Bank of St. Louis, and NBER, St. Louis, USA*
[c] *Indian Institute of Management Bangalore, Bangalore, India*

Abstract

What are the macroeconomic and distributional effects of large cash transfers when implemented at scale? We develop a two-region, small open economy heterogeneous-agent model with occupational choice, financial and real frictions, and mobility across regions. We discipline the model with baseline data from a new large-scale experiment in rural Malawi. A large-scale transfer has large, but temporary, aggregate effects on consumption and liquid assets, but a more persistent effect on productive capital and a modest impact on entrepreneurship. The consumption and entrepreneurship response is stronger for poor entrepreneurs, whereas the investment response is highest among rich entrepreneurs. The spillovers across regions are also heterogeneous: positive for workers and negative for rich entrepreneurs.

Keywords: Large-scale cash transfers, General equilibrium, Heterogeneous agents, Financial frictions, Occupational choice

1. Introduction

Large transfers to poor households, delivered either as cash or through the asset-transfer component of Graduation programs, have emerged as a promising anti-poverty policy Banerjee et al. (2015). Experimental evidence, including long-run follow-ups, shows that such interventions can generate sustained gains in consumption, assets, earnings, and reallocation into higher-return activities, though the magnitude and persistence of these effects vary across settings and program designs[1]. Yet it remains unclear how these gains scale and through which mechanisms they operate. When transfers are large enough to affect local labor supply and demand, equilibrium wage responses and spillovers to non-recipients may alter their incidence and welfare consequences. As a result, aggregate and distributional effects need not coincide with treatment effects estimated in partial-equilibrium experimental settings. What are the macroeconomic and distributional effects of large-scale cash transfers? How can experimental evidence be used to identify the mechanisms behind persistent poverty and underdevelopment? How does the source of financing shape these effects?

A new large-scale experiment in rural Malawi makes these questions especially timely. GiveDirectly is delivering lump-sum transfers of about $550 per adult under a design intended

[1] See Banerjee et al. (2015), Bandiera et al. (2017) and Banerjee et al. (2021).

to identify spillovers and broader equilibrium effects. We develop a two-region heterogeneous-agent model with occupational choice and labor markets that differ in their degree of integration, discipline it using baseline data, and generate predictions for consumption, wealth, and other variables, with particular attention to how these effects vary across households. The model provides disciplined guidance before endline outcomes are observed, helping to isolate mechanisms, trace equilibrium channels, and evaluate alternative program designs that are difficult to study from the experiment alone. As experimental evidence arrives, it will provide a sharp benchmark for assessing which margins the model captures well and where additional refinements are needed.

Our model captures several features of the general environment in which anti-poverty policies are targeted in developing countries. Households choose between wage work and entrepreneurship while facing uninsured productivity risk, stochastic access to wage employment, and tight financial constraints. They can transfer resources across time by accumulating liquid assets or productive capital, with the latter subject to partial irreversibility. This makes entrepreneurship especially risky: if a business fails, capital cannot be fully recovered. Because subsistence agriculture is quantitatively central in rural settings, we capture it as an exogenous income flow. We then embed these household decisions in a two-region small open economy with imperfect labor-market integration, so that a transfer in one region can move wages, reallocate labor, and generate spillovers in the other. This structure allows the model to speak jointly to intertemporal smoothing, occupational reallocation, local price adjustment, and spatial spillovers.

We discipline the model parameters using both the previous literature and the baseline household and enterprise data from Phase I of the GiveDirectly study in Chiradzulu, covering roughly 400 villages across the four southern Traditional Authorities. The data describe a poor, predominantly agrarian setting in which mean annual household earnings are about USD 650 and own-food production accounts for the majority of earnings on average. Market-facing entrepreneurship exists but is small-scale: about one-fifth of households operate an enterprise, most firms employ no paid workers, profits are highly skewed, and own savings are the dominant source of start-up finance. The benchmark experiment is a one-time transfer of USD 1,100, equivalent to 67% of average household consumption in the data. In the model, we represent this transfer as a one-time increase in liquid assets equal to 67% of average consumption in the stationary equilibrium.

Our benchmark experiment implies a temporary aggregate expansion in the treated region, Region 1. Households use the transfer to raise both consumption and investment, but the two margins evolve differently over time. Consumption rises immediately, whereas the capital stock builds gradually and remains elevated for longer as part of the transfer is transformed into productive business assets. Entrepreneurship rises only modestly, peaking at about 0.3 percentage points, so the aggregate response reflects limited entry and stronger expansion along the intensive margin rather than a broad wave of business creation. The expansion of entrepreneurial activity and capital raises labor demand, while some households shift out of wage work into entrepreneurship, reducing local labor supply. Together these forces push wages in Region 1 up by more than 3 percent. Temporary labor inflows from Region 2 dampen, but do not eliminate, this wage pressure. Average income rises as well, though by less than wages, and the boom eventually fades because the transfer is one-time: the extra liquidity is absorbed, the added capital depreciates, and the economy converges back to its pre-transfer stationary equilibrium.

These aggregate responses mask substantial heterogeneity within and across regions. In Region 1, the strongest consumption and entrepreneurship responses occur among lower-wealth households, for whom the transfer relaxes liquidity constraints most strongly. By contrast, the increase in productive capital is concentrated among richer incumbent entrepreneurs, who respond mainly along the intensive margin. The untreated region, Region 2, experiences a much smaller but contractionary spillover in the aggregate: wages rise modestly as labor is drawn toward Region 1, but entrepreneurship, investment, capital, and income decline. The spillover is also redistributive within Region 2. Workers and lower-wealth households benefit briefly from higher wages and a small short-run consumption gain, whereas richer entrepreneurs bear the more persistent losses through lower investment, lower wealth, and a smaller capital stock.

We also compare the benchmark economy with a fixed-price counterfactual to clarify what would be missed by focusing on partial-equilibrium treatment effects alone. Relative to a benchmark that holds prices fixed, the general-equilibrium economy generates substantially smaller increases in entrepreneurship, capital, income, and wealth because higher wages compress profits and raise the outside option of wage work. By contrast, consumption and liquid wealth differ much less across the two cases, so the main role of general equilibrium is to reshape the production response. Similarly, alternative assumptions about regional integration change the incidence of the transfer more than its aggregate welfare effect. With weaker integration, more of the transfer is absorbed into local wage growth, which benefits workers and poorer households; with stronger integration, wage pressure is smaller and more of the response appears as business expansion, which benefits entrepreneurs and wealthier households. Aggregate welfare changes very little across these cases: the main differences are distributional.

Finally, comparisons with simpler nested models sharpen the interpretation of the dynamics. Uninsured risk reduces impact spending by strengthening precautionary saving, but risk by itself generates only limited persistence. The longer-lived effects arise when the transfer can be transformed into productive capital. Comparisons with the reversible-investment case (Buera et al., 2014), the Aiyagari benchmark, and the no-shocks benchmark show that partial irreversibility is especially important for the shape of the transition: it dampens immediate entry and makes the early dynamics resemble a buffer-stock saving model, yet it also produces a long tail because those who do invest build up cautiously and unwind slowly.

1.1. Related Literature

This paper contributes first to the macro-development literature that uses quantitative general-equilibrium models to evaluate large-scale policies in developing economies. Examples include Giné and Townsend (2004) and Jeong and Townsend (2007), who study financial liberalization and financial deepening, Buera et al. (2021), who study the aggregate and distributional effects of microfinance, and Fujimoto et al. (2023), who analyze the macroeconomic effects of free secondary schooling. Most closely related to our paper is Buera et al. (2014), which studies asset grants targeting the poor in a quantitative occupational-choice model with financial frictions. Our model incorporates capital irreversibility and spatial mobility, two elements that are shown to be key to understanding the impacts. More broadly, our paper follows the agenda in Buera et al. (2023) of combining experimental evidence with structural macro frameworks to move from micro treatment effects to economy-wide policy analysis.

This paper also complements the experimental literature on large transfers and Graduation-style anti-poverty programs. Haushofer and Shapiro (2016), Banerjee et al. (2015), Bandiera et al. (2017), and Banerjee et al. (2021) document important direct and medium-run effects of such interventions on treated households, while Banerjee et al. (2022) highlights the difficulty of disentangling mechanisms within bundled programs. At the same time, impacts are not uniform across contexts: Bauchet et al. (2015) finds no lasting average effect in South India, and Barker et al. (2023) documents fading treatment effects in Ethiopia. Most closely related is Egger et al. (2022), which studies the local general-equilibrium effects of large cash transfers experimentally. Our contribution is to bring a quantitative heterogeneous-agent model to these questions, allowing us to trace equilibrium channels, study heterogeneous incidence and welfare, and evaluate counterfactual scale-ups and alternative program designs that are difficult to recover from the experiment alone.

2. Data

The data come from a cash transfer study in Chiradzulu, Malawi (Egger et al., 2024), in which GiveDirectly is implementing a universal, unconditional transfer of US$550 to each adult. Similar to Egger et al. (2022), the study design generates treatment variation across geographic units that can be used to identify both direct and spillover effects. Our analysis, however, draws only on baseline data collected prior to the transfer rollout.

Chiradzulu is a predominantly rural, agrarian district near Blantyre, with approximately 115,000 households and 450,000 residents distributed across 10 Traditional Authorities, roughly 900 villages, and 79 weekly markets. We use data from Phase I of the study, which covers the four southern Traditional Authorities, about 400 villages, combining the baseline household and enterprise modules.

Table 1 reports summary statistics from the household data. Mean annual household earnings are USD 650, or about USD 164 per capita, and households contain roughly four members on average, of whom just over two are adults. Household cash expenditure equals 1.7 times mean earnings, while total consumption, which also includes the value of produce retained for self-consumption, equals 2.52 times earnings. Gaps of this magnitude between measured consumption and measured earnings are not uncommon in rural household surveys, where income from agriculture and self-employment is difficult to capture accurately. Total income averages 1.15 times earnings and, beyond the earnings components, includes gifts and government transfers. These transfers are especially important for sustaining consumption among households with very low reported earnings.

The composition of earnings underscores the central role of subsistence production in this setting. Earnings include own food production, non-agricultural business income, labor income, crop profits, and profits from livestock activity. Own food production is by far the largest component, accounting for 63 % of earnings on average and 78 % at the median. Labor income contributes 23 %, while self-employment excluding own food production accounts for only 14 %. Although own food production is conceptually a form of self-employment, we report it separately because of its quantitative importance and because it differs from more market-oriented enterprise activity. The asset data paint a similar picture of low average wealth, substantial dispersion, and the importance of illiquid assets. Productive assets, including livestock, average only 0.08 times mean household earnings, with a median of zero, indicating that many households hold little directly productive capital. Gross assets excluding land average 0.29 times mean earnings, while gross assets including land rise to 1.56. Net worth is nearly identical at 1.55, implying that land accounts for a large share of household wealth and that debt is modest on average.[2]

Table 1: Household Summary Statistics

Variable	**Mean**	**Median**	**SD**	**N**
Aggregates				
Annual HH earnings (USD)	650.46	416.79	741.90	3,275
Income	1.16	0.80	1.20	3,275
Earnings per capita (USD)	163.99	-	-	3,275
Household Expenditure	1.70	1.08	1.95	3,275
HH members	3.97	4.00	1.71	3,275
HH adults (age $\geq$ 18)	2.18	2.00	1.00	3,257
Consumption				
Consumption	2.52	1.81	2.39	3,275
Non-durable consumption	2.50	1.80	2.37	3,275
Earnings Shares				
Self-employment share	0.14	0.00	0.29	3,046
Labor share	0.23	0.00	0.35	3,046
Own food production share	0.63	0.78	0.40	3,046
Assets				
Productive assets incl. livestock	0.08	0.00	0.24	3,275
Gross assets without land	0.29	0.11	0.46	3,275
Gross assets incl. land	1.56	0.85	2.80	3,275
Net worth	1.55	0.84	2.80	3,275

Note: Monetary variables are normalized by mean household earnings unless otherwise stated. Variables denoted in USD are converted from Malawian Kwacha at the market exchange rate prevailing during the survey period (MWK 1,724.13 = USD 1). N varies across rows due to missing values.

Table 2 partitions households into three groups motivated by the model in the next section. We classify a household as an entrepreneur household if it reports at least one active enterprise, as a worker household if it has no active enterprise but at least one member worked during the past year, and as neither otherwise. These groups account for 21.3, 54.0, and 24.6 percent of the sample, respectively. The ranking in earnings and income is clear: entrepreneur households are the most market-engaged and have the highest resources on average, worker households lie in the middle, and neither households have the lowest earnings and income. Mean household earnings are 1.34 times average sample earnings for

[2]Land is valued as the present discounted value of reported rental income using a 6% discount rate, consistent with the interest rate in the model.

entrepreneur households, compared with 1.00 for worker households and 0.71 for neither households; the same ordering appears for total income. Consumption is also highest for entrepreneur households, though the gap between worker and neither households is much smaller than the gap in earnings. The number of adults is very similar across groups, so these differences primarily reflect economic activity rather than household composition.

Table 2: Group Wise: Summary Statistics

Variable	Mean	Median	SD	N
Panel A: Ent (share = 21.32%)				
Aggregates				
HH earnings	1.34	1.01	1.24	696
Income	1.49	1.13	1.27	696
HH adults (age $\geq$ 18)	2.28	2.00	1.00	692
Consumption				
Consumption	3.02	2.27	2.56	696
Shares				
Own food production share	0.46	0.45	0.37	689
Self-employment share	0.43	0.36	0.36	689
Labor share	0.11	0.00	0.24	689
Panel B: Worker (share = 54.02%)				
Aggregates				
HH earnings	1.00	0.65	1.13	1,751
Income	1.16	0.80	1.19	1,751
HH adults (age $\geq$ 18)	2.18	2.00	0.97	1,742
Consumption				
Consumption	2.39	1.75	2.26	1,751
Shares				
Own food production share	0.59	0.69	0.40	1,662
Self-employment share	0.04	0.00	0.16	1,662
Labor share	0.37	0.22	0.39	1,662
Panel C: Neither (share = 24.56%)				
Aggregates				
HH earnings	0.71	0.27	1.00	825
Income	0.87	0.49	1.09	825
HH adults (age $\geq$ 18)	2.10	2.00	1.05	820
Consumption				
Consumption	2.37	1.60	2.48	825
Shares				
Own food production share	0.90	1.00	0.26	692
Self-employment share	0.10	0.00	0.26	692
Labor share	0.00	0.00	0.00	692

Notes: Monetary variables are normalized by mean household earnings = USD 650.46. N varies across rows due to missing values.

The composition of earnings also lines up closely with the group definitions. Entrepreneur households derive 43% of earnings from self-employment and 46% from own food production, while labor income contributes only 11%. Worker households rely much more heavily on labor income, which accounts for 37% of earnings, but own food production still represents the largest single component at 59%; self-employment outside own food production is negligible at 4%. The neither group is essentially detached from market-facing activity; labor income is zero by construction, self-employment is very limited, and 90% of earnings comes from own food production, with a median share of 100%.

Notably, subsistence agriculture is not confined to the neither category but remains important across all three groups. Our model is designed to capture the margins most relevant for the first two groups, namely households that engage in market-facing activity through enterprise operation or labor supply. For this reason, the moments used in calibration rely primarily on entrepreneur and worker households. At the same time, because subsistence production is quantitatively central in this setting but not the focus of our mechanism, we treat it in the model as an exogenous income flow rather than an endogenous occupational margin.

Table 3: Enterprise Statistics

Variable	**Mean**	**Median**	**SD**	**Max**	**N**
Annual profit (USD)	999.62	389.76	1837.11	10440.05	2,336
Value added	−0.55	−0.31	8.69	31.82	266
Productive assets excl. land & buildings	0.11	0.00	0.39	3.02	2,406
Productive assets incl. land & buildings	1.73	0.02	3.74	23.56	2,340
Own savings (dummy)	0.75	1.00	0.44	-	2,219
Paid employees (lower bound, capped at 3)	0.23	0.00	0.60	-	2,406
Business age (years)	8.63	5.31	9.16	36.9	2,309

Notes: Monetary variables are normalised by mean annual profit unless otherwise stated. Variables denoted in USD are converted from Malawian Kwacha at the market exchange rate prevailing during the survey period (MWK 1,724.13 = USD 1). **Max** refers to the maximum value in the data after winsorization at the 99th percentile. N varies across rows due to missing values.

Table 3 reports summary statistics for active enterprises. Enterprise profits are sizable but highly skewed. Mean annual profit is about USD 1,000, while the median is only USD 390 and the standard deviation exceeds USD 1,800, indicating a long right tail. At the same time, these are very small businesses by standard measures. The median enterprise employs no paid workers, and the mean number of paid employees is only 0.23 even when employment is capped at three. Additionally, roughly 85 percent of firms have no paid employees at all, so most enterprises in the sample are effectively self-operated. Sector composition is also concentrated in low-scale activities: about half of firms are food vendors, underscoring that the relevant enterprise margin in this setting is microenterprise rather than larger-scale formal business activity. Businesses are also not especially young, with a median age of just over five years and a mean of 8.6 years.

Asset holdings are highly uneven and concentrated in illiquid forms. Excluding land and buildings, productive assets average only 0.11 times mean annual profit and the median is zero, indicating that many firms operate with almost no movable productive capital.

Including land and buildings raises average productive assets to 1.73 times mean annual profit, while the median remains only 0.02, revealing a highly skewed distribution driven by a relatively small number of firms with substantial real assets. As in the household balance-sheet measures, land and buildings are valued using reported rental values and a 6 percent discount rate, consistent with the interest rate in the model. Financing is also heavily internal: 75 percent of entrepreneurs report using their own savings to start the business, consistent with limited access to external credit. We also construct a measure of value added as revenues minus input costs, but this variable is available for only a small subset of firms and is extremely noisy. As emphasized by de Mel et al. (2009), measuring profits and value added in enterprise surveys is difficult, and that is true in our setting as well. For this reason, reported profits are more informative than value added for describing enterprise heterogeneity and guiding calibration.

3. Model

Time is continuous. An infinitely lived agent chooses consumption c_t, investment i_t, and whether to operate a business, $e_t \in \{0, 1\}$, where $e_t = 1$ denotes entrepreneurship and $e_t = 0$ denotes non-entrepreneurship.

The agent can save in two assets. The first is a liquid asset b_t, which can be adjusted freely and earns interest rate r. The second is illiquid productive capital k_t, which evolves through investment and depreciation and is subject to partial irreversibility through the adjustment-cost function $\Phi(i_t, k_t)$.

A key feature of the environment is that wage employment is not always available. We capture this with an indicator $x_t \in \{0, 1\}$: when $x_t = 1$, the agent has access to the labor market and earns wage w_t, provided he does not operate a business; when $x_t = 0$, wage work is unavailable. A second feature of rural economies is dependence on subsistence agriculture, which we model as an exogenous income flow y_t^s.

In addition, the agent faces idiosyncratic entrepreneurial productivity risk z_t. The processes for z_t and x_t evolve stochastically over time. These risks are uninsured, and the agent is also subject to credit constraints. Formally, the agent solves

$$\max_{\{c_t, i_t, e_t\}_{t \geq 0}} \quad \mathbb{E}_0 \int_0^\infty e^{-\rho t} \frac{c_t^{1-\sigma}}{1-\sigma} \, dt \tag{1}$$

subject to

$$db_t = \Big[y_t^s + (1 - e_t) x_t w_t + e_t \Pi(z_t, k_t) + r b_t - c_t - i_t - \Phi(i_t, k_t) \Big] dt, \tag{2}$$

$$dk_t = (i_t - \delta k_t)\, dt, \tag{3}$$

$$z_t \in \{z_1, z_2, z_3, z_4\}, \tag{4}$$

$$x_t \in \{0, 1\}, \qquad e_t \in \{0, 1\}, \tag{5}$$

$$b_t \geq 0, \qquad k_t \geq 0. \tag{6}$$

When the agent chooses entrepreneurship ($e_t = 1$), he operates a technology with idiosyncratic productivity z_t and capital k_t, and hires labor l_t at wage w_t. Profits are therefore given by

$$\Pi(z_t, k_t) = \max_{l_t} \left\{ z_t \left(k_t^{\alpha} l_t^{1-\alpha}\right)^{1-\nu} - w_t l_t \right\}. \tag{7}$$

3.1. Regions

We consider an economy with two ex ante identical regions, 1 and 2, each with population share 1/2. A one-time cash transfer is implemented in region 1. Labor markets are only partially integrated: residents may temporarily supply labor in the other region when wage differences make this profitable.

In the initial stationary equilibrium, the two regions are symmetric. Wages, entrepreneurial returns, and occupational composition are identical across regions, so no individual has an incentive to move.

The transfer to region 1 temporarily breaks this symmetry. By raising local labor demand, it may increase wages in region 1 relative to region 2. This creates an incentive for some workers from region 2 to temporarily work in region 1. We assume that temporary movers can work only as wage workers at the destination: they earn the destination wage $w_{1,t}$, but cannot operate an entrepreneurial activity there. This captures the idea that entrepreneurship requires location-specific knowledge, networks, or capital that do not easily move across regions, at least in response to temporary shocks.

Migration is modeled as a reduced-form earnings wedge. Each worker in region 2 is characterized by an idiosyncratic migration cost factor $c \geq 1$, drawn independently of all other characteristics from a Pareto distribution with shape parameter $\theta > 0$:

$$F(c) = 1 - c^{-\theta}, \qquad c \geq 1. \tag{8}$$

The factor c summarizes the frictions associated with temporary movement, including travel costs, the disutility of being away from home, and temporary accommodation costs. A worker with cost factor c who works in region 1 receives effective labor income $w_{1,t}/c$. When $c = 1$, the worker faces no mobility friction; when $c > 1$, migration reduces the effective return from destination work.

Residents first choose occupation in their home region. Conditional on being wage workers, region-2 workers then decide whether to supply labor locally or temporarily in region 1. A worker in region 2 earning the local wage $w_{2,t}$ chooses temporary work in region 1 whenever

$$\frac{w_{1,t}}{c} \geq w_{2,t}, \tag{9}$$

or equivalently,

$$c \leq \frac{w_{1,t}}{w_{2,t}}. \tag{10}$$

Since all workers in region 2 face the same local and destination wages, the cutoff $w_{1,t}/w_{2,t}$ is common across them. The fraction of region-2 workers who temporarily move is therefore given by the Pareto CDF evaluated at this threshold.

Let $N^w_{2,t}$ denote the mass of residents of region 2 who choose wage work at time t. This mass is determined by occupational choice within region 2 and may vary over time as local returns change. Temporary movement does not alter residency: workers who temporarily work in region 1 remain residents of region 2 and are still counted in $N^w_{2,t}$.

Let M_t denote the mass of region-2 residents who are temporarily supplying labor in region 1 at time t. Then

$$M_t = \left[1 - \left(\frac{w_{1,t}}{w_{2,t}}\right)^{-\theta}\right] N^w_{2,t} \qquad \text{if } w_{1,t} > w_{2,t}, \tag{11}$$

and $M_t = 0$ otherwise. The effective labor supplied in each region is therefore

$$L_{1,t} = N^w_{1,t} + M_t, \tag{12}$$

$$L_{2,t} = N^w_{2,t} - M_t. \tag{13}$$

In the initial stationary equilibrium, $w_{1,t} = w_{2,t}$, so the threshold equals one and $M_t = 0$. No temporary movement occurs.

The parameter θ governs the degree of labor market integration across regions. Intuitively, θ determines how responsive temporary migration is to a given wage gap. Low values of θ imply that most workers face very large effective mobility costs, so even sizable wage differences induce little movement. High values of θ imply that mobility costs are concentrated near one, so even small wage differences generate substantial labor reallocation. In this sense, θ provides a simple reduced-form measure of how integrated the two regional labor markets are.

This formulation nests three benchmark cases that will be useful throughout the analysis. First, as $\theta \to 0$, mobility becomes negligible and the two regions behave as effectively segmented labor markets. Second, for intermediate values of θ, only a subset of workers respond to wage differences, so labor markets are partially integrated. Third, as $\theta \to \infty$, any positive wage gap induces all workers to move, so the economy converges to the case of full labor market integration.

The temporary nature of movement is central. Workers re-evaluate the location of labor supply at each instant and move only when the contemporaneous wage gap makes it profitable. Because the transfer is one-time, the economy converges to its original symmetric stationary equilibrium over time. Temporary movement therefore shrinks back to zero and the economy converges to its original symmetric stationary equilibrium.

4. Calibration

This section describes how we discipline the model using micro- and macro-level data and assess its empirical relevance. We distinguish between parameters assigned externally, based on standard values from the literature due to the lack of available data, and those calibrated internally to match moments from the Malawian data.

4.1. Externally assigned parameters.

The coefficient of relative risk aversion is set to $\sigma = 1.5$, and the depreciation rate to $\delta = 0.06$, values that are widely used in related quantitative work. We treat the economy as small and open, so the interest rate is taken as exogenous; in the benchmark specification, we set $r = 0$ to capture the lack of saving opportunities (Dupas and Robinson, 2013). For the entrepreneurial technology, we choose $\alpha = 0.33$ and $\nu = 0.2$, which lie in the range commonly used in the literature.

The adjustment cost function Φ takes the form:

$$\Phi(i,k) = \begin{cases} \frac{\kappa}{2}\left[\frac{i}{k+\bar{k}}\right]^2 (k+\bar{k}) & i \geq 0 \\ -(1-\phi)i + \frac{\kappa}{2}\left[\frac{i}{k+\bar{k}}\right]^2 (k+\bar{k}) & i < 0 \end{cases}$$

The term $-(1-\phi)i$ for $i < 0$ introduces partial irreversibility: when an entrepreneur disinvests, she recovers only $\phi \leq 1$ per unit of capital sold. The quadratic term, governed by κ and the small constant $\bar{k}$ are used to ensures smoothness of the problem. We choose $\phi = 0.35$, a relatively low value in the literature (Kermani and Ma, 2023), but present results for the reversible case, which encompasses higher values (Baley and Blanco, 2026). We set $\kappa = \bar{k} = 0.1$.

To parameterize the labor-market opportunity process, we note that wage employment in Malawi is closely tied to the agricultural season. In a typical year, the main period of labor demand spans roughly October/November through April/May, reflecting labor needs during land preparation, planting, and weeding, with some additional demand during harvest. This is consistent with our data, in which workers report being employed for an average of 5.5 months per year. We therefore set the transition rates governing labor-market access to

$$\left(\lambda_x^{0\to1},\, \lambda_x^{1\to0}\right) = \left(\tfrac{12}{6},\, \tfrac{12}{6}\right).$$

The parameter θ governs the responsiveness of temporary migration to wage differences. We calibrate it from Bryan et al. (2014), who study seasonal migration in rural Bangladesh. In their control group, 36.0% of eligible households send at least one seasonal migrant, and the paper reports a descriptive earnings comparison of roughly 100 Tk/day for migrants at destination versus 65 Tk/day for a sub-sample of non-migrants in salaried work at origin. Imposing these moments on our migration rule yields $\theta \approx 1.04$, which we take as a reduced-form benchmark for migration responsiveness.

4.2. Internally calibrated parameters

We calibrate remaining parameters by matching model-generated moments to their empirical counterparts. The subsistence flow y^s is pinned down by targeting the ratio of aggregate own-food consumption to aggregate earnings, which equals 57.4% in the data. The discount rate is chosen so that the stationary equilibrium reproduces the observed aggregate net-worth to earnings ratio of 1.31.

We model entrepreneurial productivity in logs, $\tilde{z}_t \equiv \log z_t$, as a four-state continuous-time Markov chain. The chain is designed to provide a parsimonious approximation to a latent AR(1) process in logs,

$$\tilde{z}_{t+1} = \rho_z \tilde{z}_t + \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim \mathcal{N}(0, \sigma_\varepsilon^2). \tag{14}$$

Table 4: Externally Set Parameters

Parameter	**Description**	**Value**
σ	Risk aversion	1.5
δ	Depreciation rate	0.06
r	Interest rate	0.00
ϕ	Partial irreversibility	0.35
$\kappa, \bar{k}$	Quadratic adjustment cost	$0.1, 0.1$
α	Production function (Capital)	0.33
ν	Span of control	0.2
$\lambda_x^{0\to1}, \lambda_x^{1\to0}$	Wage opportunity process	$2, 2$
θ	Migration parameter	1.04

Table 5: Internally Calibrated Parameters

Parameter	Description	Value	Target
y^s	Subsistence flow	1.09	Agg own food / Agg earnings
ρ	Discount rate	0.046	Agg net worth / Agg earnings
$(\rho_z, \sigma_\varepsilon)$	Productivity process	(0.85, 0.26)	Literature, Entrepreneurs mean/median profits

As a discretization choice, we fix the invariant mass vector ex ante at

$$p = (0.25, 0.25, 0.25, 0.25),$$

so that each state carries one quarter of the stationary mass. Given $(\rho_z, \sigma_\varepsilon)$, we assign one representative grid point to each quartile and scale the resulting support to match the stationary dispersion implied by the AR(1). We then normalize productivity levels so that $\mathbb{E}[z] = 1$. Finally, for this fixed support and invariant mass vector, we recover the generator λ_z so that the chain has stationary distribution p and reproduces annual persistence ρ_z. Thus, in the discretization, σ_ε governs the spread of the productivity grid, while ρ_z governs the persistence of movements across states.

Because our data are cross-sectional, they are uninformative about persistence. We therefore set $\rho_z = 0.85$ following the external evidence in Asker et al. (2014), and choose σ_ε to match the mean-to-median ratio of entrepreneurial profits, which equals 2.56 in the data.

5. Benchmark Results

5.1. Benchmark Experiment

As noted earlier, the cash transfer is $550 per adult. With roughly 2 adults in the average household, this amounts to about $1,100 per household. In the baseline data, this is about 67% of mean household consumption and roughly equal to mean household cash expenditure. While the transfer is also large relative to reported household earnings ($\times 1.69$), income is measured with considerable noise in this setting, so we view consumption as more reliable benchmarks for scaling the intervention. Our model is formulated at the household level and abstracts from children and intra-household allocation, so we map the adult-level intervention into a single lump-sum household transfer. We set the transfer (Tr) equal to 67% of mean household consumption in the stationary equilibrium, matching the empirical ratio. Formally, the transfer is modeled as a one-time, unanticipated increase in the liquid assets. Each agent's state transitions as

$$(b, k, z_i, x_i) \longrightarrow (b + Tr, k, z_i, x_i)$$

where Tr denotes the transfer amount.

5.2. Aggregate Implications

Figure 1 shows the response of Region 1 aggregates to the transfer. Liquid wealth rises on impact by roughly 63 percent of steady-state income. In the aggregate, households use this windfall both to raise current consumption and to finance a sharp burst of investment: consumption increases by about 17 percent on impact, while investment jumps by a similar amount relative to steady-state income. Liquid wealth is then depleted rapidly, but total wealth declines more slowly because part of the transfer is converted into productive capital.

The capital stock therefore rises with a lag and peaks only after several years, while entrepreneurship increases very modestly, peaking at 0.3 pp. This pattern is consistent with the transfer primarily relaxing liquidity constraints and allowing existing businesses to scale up, with more limited movement along the entry margin. The expansion of entrepreneurial activity and capital raises labor demand, and the shift of some households into entrepreneurship reduces local wage-labor supply. Together, these forces push wages in Region 1 up by more than 3 percent over time. Because Region 1 is only partially integrated with the untreated Region 2, the wage increase induces net in-migration, which dampens but does not fully offset the local labor-market tightening. Average income rises as well, though by less than wages, and remains above baseline even after liquid wealth has largely disappeared. Because the transfer is one-off, it does not alter long-run fundamentals. As the extra liquidity is consumed and invested, the added capital depreciates, and idiosyncratic shocks reshuffle households across the wealth distribution, Region 1 converges back to its pre-transfer stationary equilibrium.

5.3. Micro-Implications

Figure 2 examines the micro-level impact across workers and entrepreneurs within each wealth quartile, showing both what drives the aggregate response and how the transfer has heterogeneous effects across households. Among low- and middle-wealth households, the main effects are a large consumption response and a sizeable increase in entrepreneurship, consistent with the transfer relaxing liquidity constraints and enabling entry among marginal businesses. Among high-wealth households, and especially baseline entrepreneurs, the response is instead along the intensive margin: capital rises immediately and strongly, while the entrepreneurship response is negligible. The aggregate increase in entrepreneurship is therefore driven primarily by lower-wealth households, whereas the aggregate increase in productive capital is driven disproportionately by richer incumbent entrepreneurs.

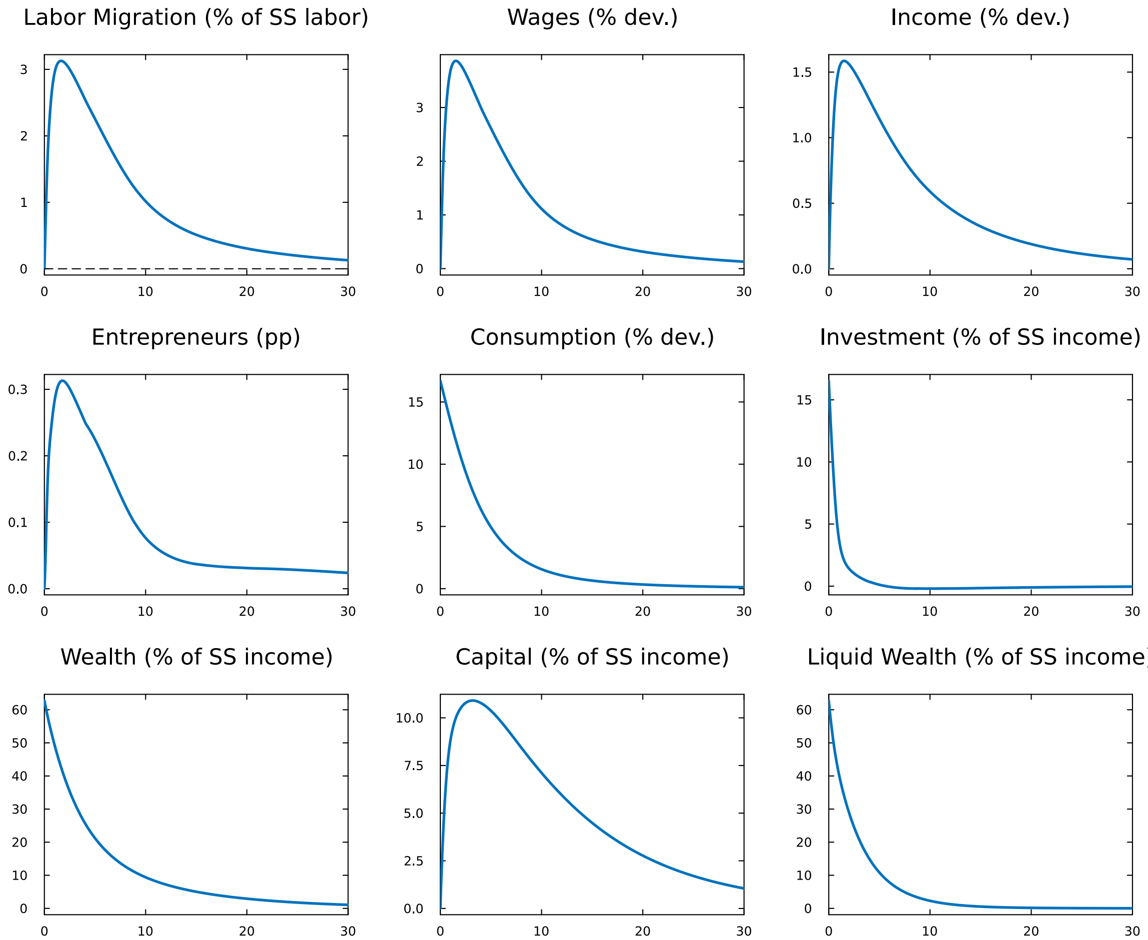

Figure 1: The figure shows the dynamic responses of the main aggregate variables in Region 1, the cash-transfer region, following the transfer. Region 1 is partially integrated with the rest of the economy. Outcomes are reported as percent deviations from their stationary equilibrium values. For selected variables, time paths are as a share of aggregate steady-state income.

5.4. Spatial Spillovers

Figure 3 shows that the untreated Region 2 experiences a contractionary spillover from the transfer to Region 1. Higher wages in Region 1 induce temporary out-migration from Region 2, and the resulting fall in local labor supply raises wages in Region 2 by almost 0.8 percent. Because Region 2 receives no direct transfer, its entrepreneurial activity contracts: entrepreneurship declines by up to 0.12 pp, investment turns negative, and the capital stock falls over time. Average income in Region 2 therefore declines even as wages rise.

The joint behavior of consumption, total wealth, and liquid wealth indicates that households in Region 2 initially respond by reallocating their portfolios away from business capital. Total wealth falls throughout and the capital stock drops, but liquid wealth rises temporarily. This pattern is consistent with households cutting investment and allowing business capital to run down as local entrepreneurial opportunities worsen. The short-run increase in liquid resources supports a very small increase in consumption on impact, despite the decline in income. Over time this liquidity buffer fades: as the smaller capital stock depresses earn-

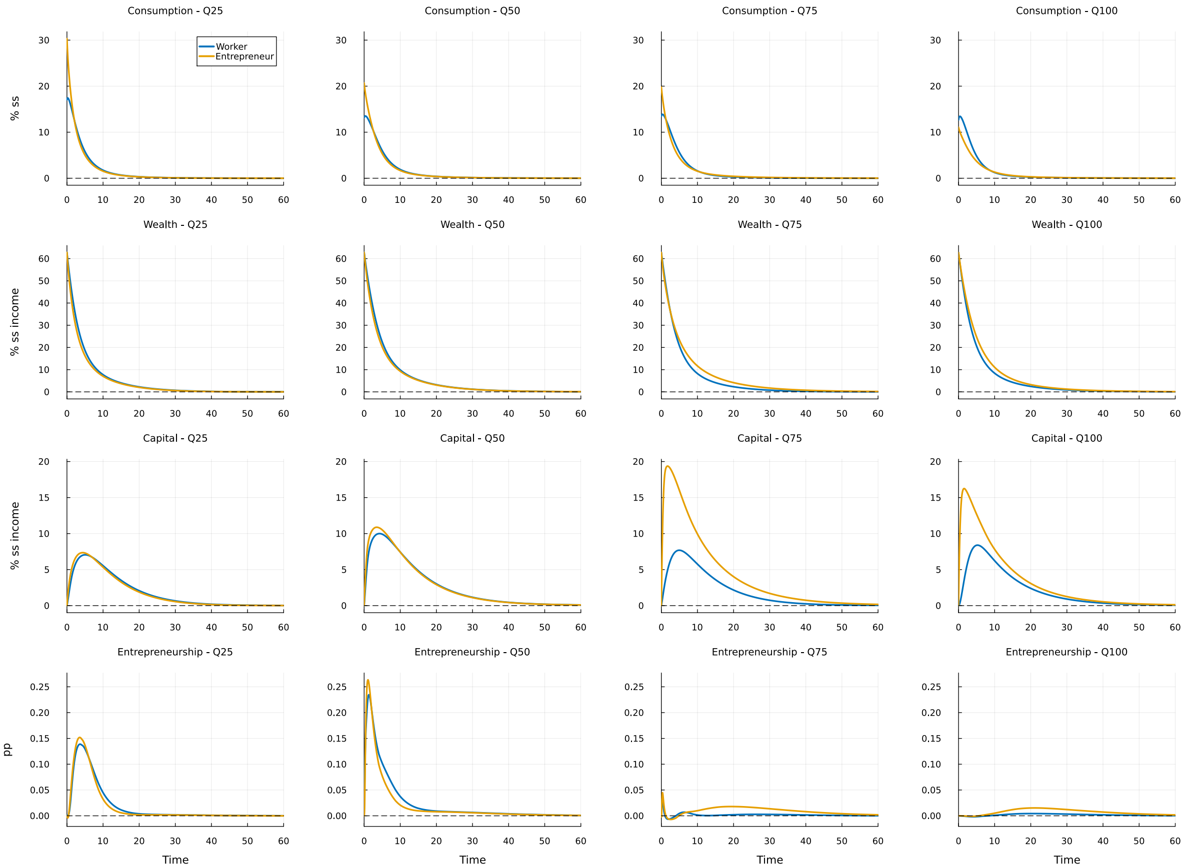

Figure 2: The figure shows the dynamic responses in Region 1 by baseline wealth quartile and occupation. The blue line denotes workers and the yellow line denotes entrepreneurs within each quartile. The plotted responses compare outcomes in the cash-transfer region (Region 1) with those for the same baseline group in the no-transfer counterfactual.

ings and liquid balances are exhausted, consumption falls below baseline and only gradually recovers.

Figure 4 shows that the aggregate effects in Region 2 are highly heterogeneous across households. The short-run consumption gain is concentrated among workers and lower-wealth households, who benefit most directly from the wage increase. By contrast, the largest declines in wealth and capital occur among richer baseline entrepreneurs, especially in the top two wealth quartiles. For these households, the response is along the intensive margin: they remain entrepreneurs but sharply reduce the scale of their businesses. The negative entrepreneurship response is more pronounced in the lower and middle parts of the wealth distribution, suggesting that marginal businesses are the most likely to exit when Region 2 becomes a less attractive place to operate.

Taken together, these patterns imply that the cross-region spillover is redistributive within Region 2. Workers and poorer households benefit temporarily from higher wages, while incumbent entrepreneurs bear most of the persistent losses through lower investment and a smaller capital stock. This heterogeneity reconciles the aggregate results: the modest positive consumption response on impact is driven by wage gains and a temporary increase in liquid assets, whereas the sustained declines in income, wealth, and capital reflect the

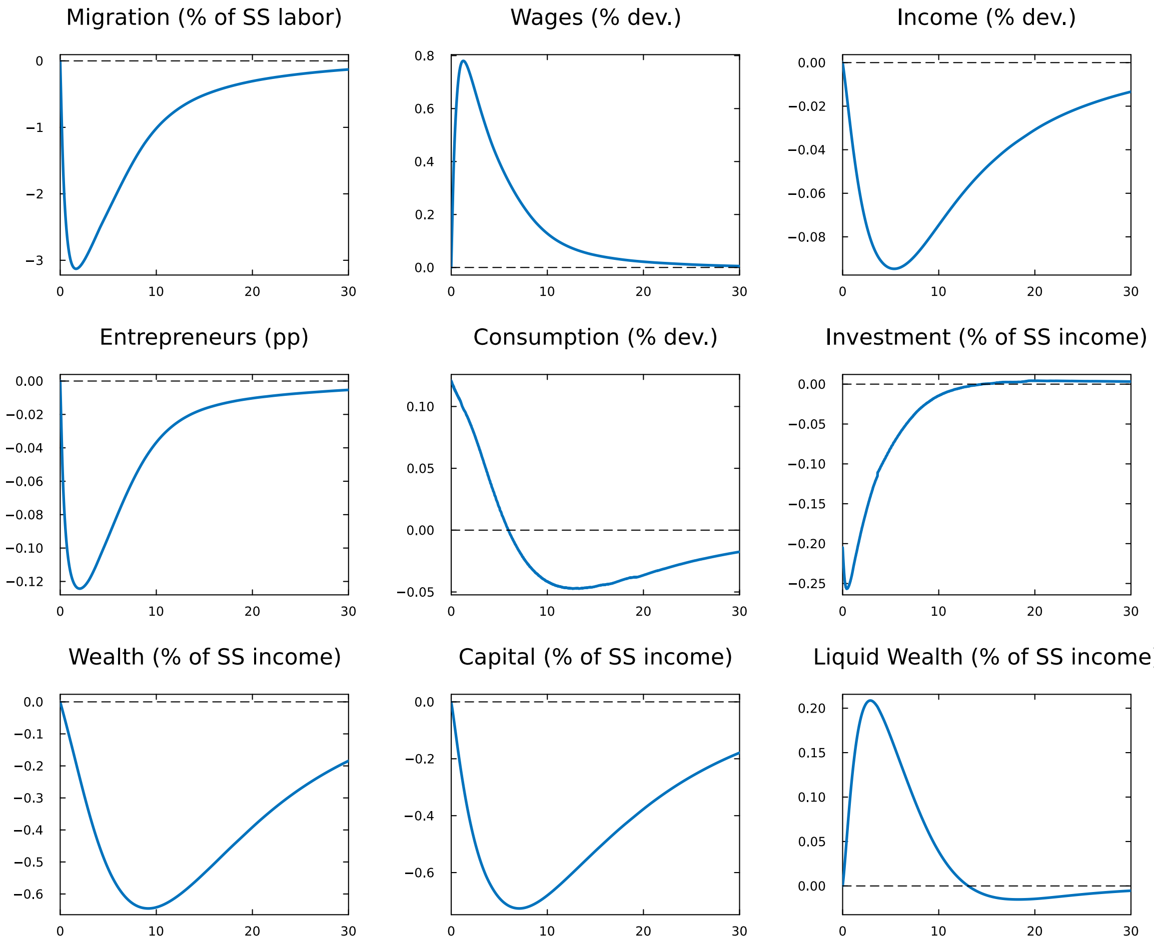

Figure 3: The figure shows the dynamic responses of the main aggregate variables in Region 2 following the cash transfer in Region 1.Outcomes are reported as percent deviations from their stationary equilibrium values. For selected variables, time paths are as a share of aggregate steady-state income..

contraction of the entrepreneurial sector.

6. Welfare, Price Effects, and Spatial Integration

We summarize welfare using a consumption-equivalent variation (CEV). Let $V_{ss}(s)$ denote the household value function in the pre-transfer stationary equilibrium, and let G_{ss} denote the associated stationary distribution over states s. Let G_{Tr} denote the distribution immediately after the surprise transfer, that is, after liquid wealth has been shifted up by the transfer but before the subsequent transition has unfolded. Under scenario $i \in \{\text{Fixed Prices}, \text{Zero}, \text{Benchmark}, \text{Full}\}$, let $V_0^i(s)$ denote the value at date 0 for a household that starts from post-transfer state s and then faces the entire future transition implied by scenario i. The aggregate consumption-equivalent welfare gain is

$$\omega_i = \left(\frac{\int V_0^i(s)\, dG_{Tr}(s)}{\int V_{ss}(s)\, dG_{ss}(s)} \right)^{\frac{1}{1-\sigma}} - 1.$$

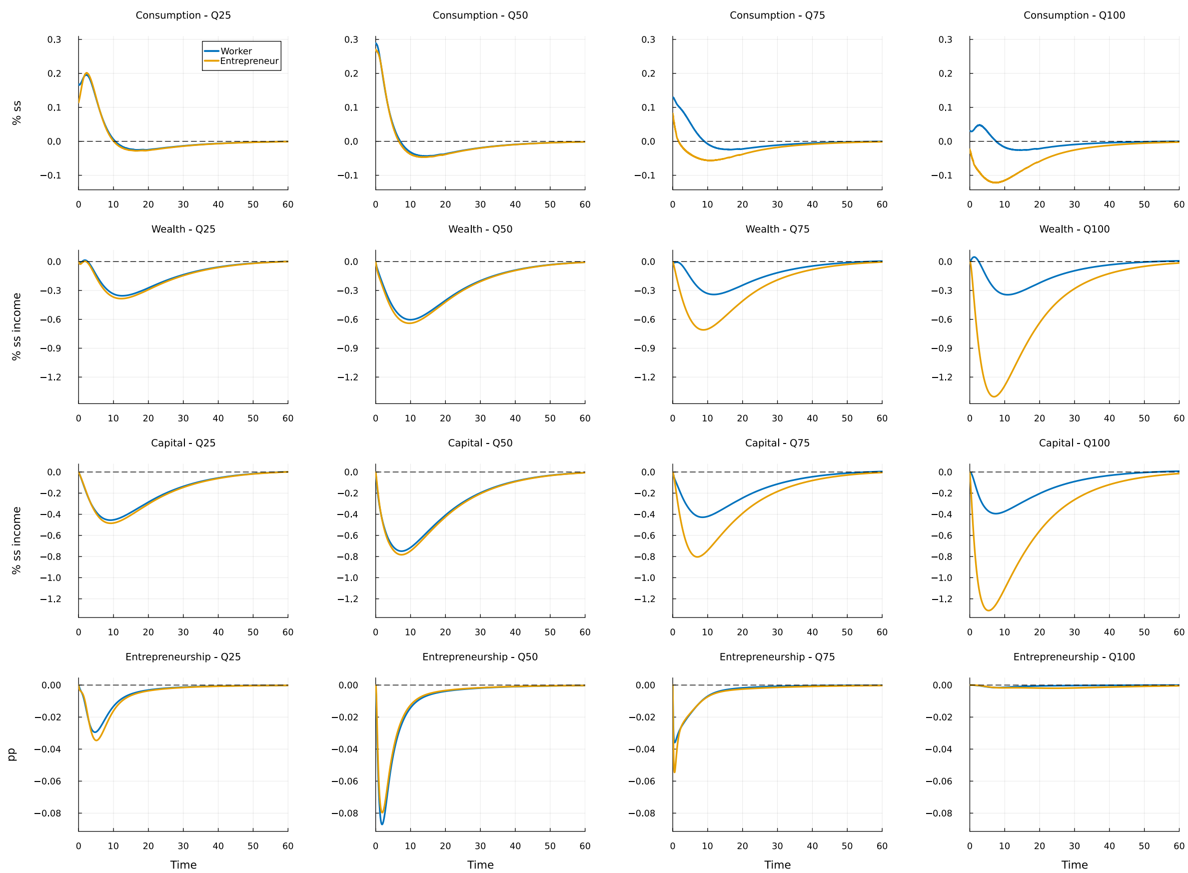

Figure 4: The figure shows the dynamic responses in Region 2 by baseline wealth quartile and occupation. Blue line denotes workers and the yellow line denotes entrepreneurs from that quartile. The plotted responses compare outcomes in the spillover economy (Region 2) induced by the transfer to Region 1 to outcomes for the same baseline group in the no-transfer counterfactual.

This measure is the uniform proportional increase in consumption, relative to the stationary benchmark, that makes households indifferent between remaining in the pre-transfer stationary equilibrium and receiving the transfer under scenario i. In the fixed-prices case, $V_0^i(\cdot) = V_{ss}(\cdot)$ as a function of the state, so welfare differs from the benchmark only because the transfer shifts the distribution from G_{ss} to G_{Tr}; in the remaining scenarios, V_0^i also incorporates the endogenous price path generated by the transfer.

To study heterogeneity, we compute the same object within baseline groups g, such as workers versus entrepreneurs or baseline wealth quartiles. Let G_{ss}^g denote the stationary distribution restricted to households in group g, and let G_{Tr}^g denote the post-transfer distribution for those same households. We then define

$$\omega_i(g) = \left(\frac{\int V_0^i(s)\, dG_{Tr}^g(s)}{\int V_{ss}(s)\, dG_{ss}^g(s)} \right)^{\frac{1}{1-\sigma}} - 1.$$

Thus, $\omega_i(g)$ is the average consumption-equivalent welfare gain for households that belonged to group g before the transfer arrived.

Table 6: Welfare Gains: Region 1 (%)

		Integration		
	Fixed Prices	Zero	**Benchmark**	Full
Aggregate	2.80	2.78	2.78	2.77
Baseline occupation				
Workers	2.82	2.91	2.89	2.84
Entrepreneurs	2.79	2.69	2.70	2.71
Baseline wealth quartile				
Bottom 25%	3.13	3.31	3.27	3.19
25–50%	2.91	2.99	2.97	2.93
50–75%	2.78	2.71	2.71	2.72
Top 25%	2.37	2.09	2.13	2.20

6.1. Partial vs. General Equilibrium

Figure 5 compares the baseline response in Region 1 to a partial-equilibrium counterfactual in which local prices are held fixed at their pre-transfer levels. Liquid wealth paths are nearly indistinguishable across the two cases, and consumption paths are also close. The differences emerge on the production side: with prices fixed, the transfer induces a substantially larger increase in entrepreneurship, a stronger and more persistent buildup of capital, and larger gains in income and wealth. The fixed-price case also generates a large excess demand for labor—latent labor-market pressure that, in the benchmark equilibrium, is absorbed by higher wages and labor inflows.

This comparison reveals that local general-equilibrium price effects substantially attenuate the productive response to the transfer. When prices are fixed, firms expand without facing higher labor costs, and households can shift into entrepreneurship without a rise in the wage outside option. In the benchmark equilibrium, by contrast, the rise in labor demand pushes Region 1 wages up, which simultaneously compresses business profits and makes wage work more attractive for marginal households. This feedback reduces entry, slows capital accumulation, and lowers the resulting gains in income and wealth. A fixed-price exercise therefore overstates the medium-run production impact of the transfer.

These quantity differences, however, do not translate into large differences in aggregate welfare. Welfare gains are nearly identical across the two cases—2.80 percent under fixed prices versus 2.78 percent in the benchmark—because endogenous price adjustment is redistributive rather than aggregate. Fixed prices favor entrepreneurs and wealthier households, who benefit from stronger business expansion and cheaper labor: welfare gains are higher for baseline entrepreneurs (2.79 versus 2.70) and for the top half of the wealth distribution. The benchmark equilibrium instead favors workers and poorer households, who gain directly from higher local wages: welfare gains are higher for baseline workers (2.89 versus 2.82) and for the bottom half of the wealth distribution, especially the bottom quartile (3.27 versus 3.13). The wage response thus dampens the production side of the local boom while shifting gains toward households that rely more on labor income and hold less wealth.

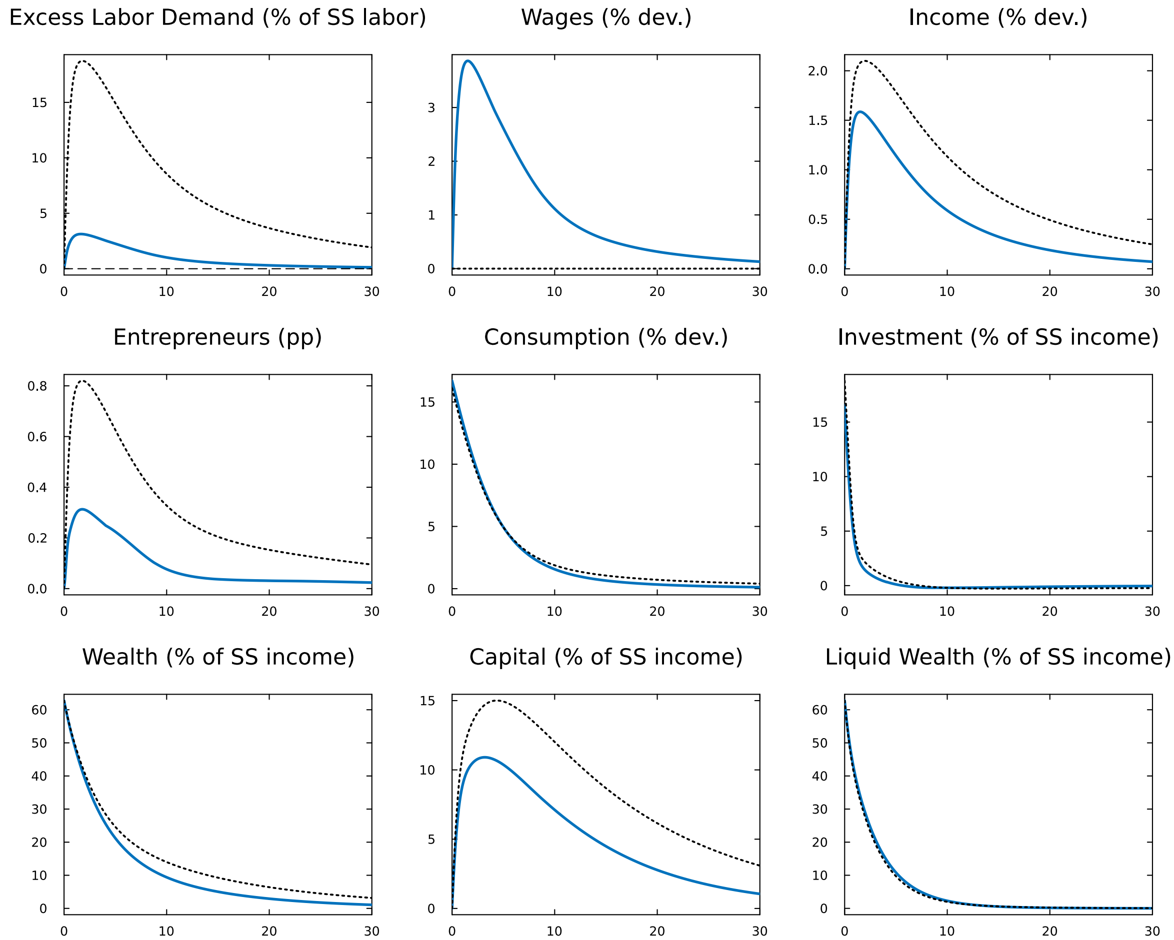

Figure 5: The figure compares the dynamic responses of the main aggregate variables in Region 1 under partial integration with those under fixed prices. The solid blue line shows the general-equilibrium path with labor-market clearing, while the black dotted line shows the fixed-price benchmark with no wage adjustment.

6.2. Alternative Spatial Spillovers

Figure 6 compares Region 1 under three labor-market regimes: full segmentation (zero integration), the benchmark of partial integration, and full integration. Consumption and liquid wealth move very similarly on impact across the three cases, since households receive the same transfer and initially split it between spending and saving in much the same way. The differences emerge in the general-equilibrium response of the local labor market: as integration rises, the wage response becomes smaller, while entrepreneurship, capital, income, and total wealth become larger and more persistent. Segmentation produces the sharpest wage increase but the weakest business expansion; full integration dampens the wage response the most and generates the strongest business expansion; the benchmark lies in between.

The mechanism runs through the elasticity of local labor supply. When Region 1 is segmented, the transfer-induced increase in labor demand cannot be met by inflows from Region 2, so local labor scarcity pushes wages up sharply. Those higher wages raise labor costs for firms and increase the wage outside option for marginal households, which discourages entry and limits capital accumulation. More of the transfer is therefore absorbed as

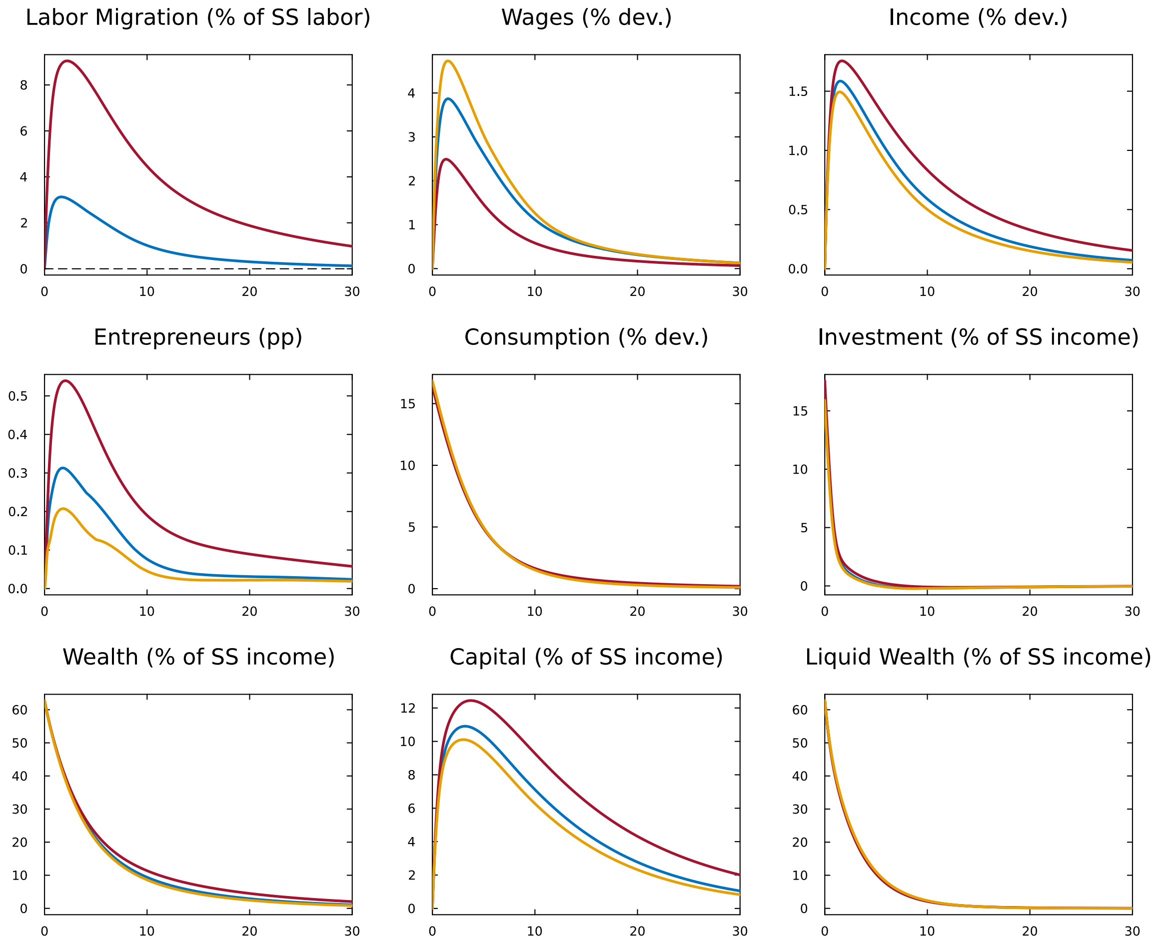

Figure 6: The figure compares the dynamic responses of the main aggregate variables in Region 1 under different labor-market integration assumptions. The solid blue line corresponds to partial integration, the red line to full integration, and the yellow line to complete segmentation.

higher wages and less is translated into productive expansion. Under full integration the opposite holds: labor can reallocate across regions, wage pressure is weaker, and firms expand at lower cost, so more of the transfer shows up in entrepreneurship, capital, and income rather than in local wages. The benchmark displays the intermediate pattern one would expect from a partially elastic labor supply.

As in the fixed-price comparison, these quantity differences mostly reflect redistribution rather than aggregate welfare gains. Aggregate welfare is almost identical across the three cases—2.78 percent under zero integration, 2.78 percent in the benchmark, and 2.77 percent under full integration—but the incidence shifts systematically with the degree of integration. Full segmentation benefits workers and poorer households the most, since it generates the sharpest local wage increase: baseline workers gain 2.91 percent under zero integration versus 2.89 in the benchmark and 2.84 under full integration, and the bottom quartile gains 3.31 percent versus 3.27 and 3.19. Full integration instead favors entrepreneurs and wealthier households, who benefit from lower labor costs and stronger business expansion: baseline entrepreneur welfare rises from 2.69 to 2.70 to 2.71 across the three cases, and the top quartile's gain rises from 2.09 to 2.13 to 2.20. Stronger integration thus leaves the aggregate welfare effect of the transfer nearly unchanged but shifts the gains away from workers and toward incumbent businesses and wealthier households.

7. Model Comparisons

This section studies the role of the different forces operating in our benchmark model, which we refer to as the *Main Model*. It features entrepreneurial productivity risk z_t, labor-market risk x_t, endogenous investment, and partial irreversibility. We consider a series of nested specifications that shut these forces down one at a time and compare their aggregate responses. In every case, the experiment is a one-time transfer that raises initial liquid wealth by 62.8% of the model's own stationary gross output, matching the transfer size in the benchmark.

Reversible: Uncertainty + Investment, No Irreversibility

To shut down the role of irreversibility, we let entrepreneurs choose productive capital statically each period, subject only to the borrowing constraint. This is similar to Buera et al. (2014). Formally, the household solves

$$\max_{\{c_t,e_t\}_{t\geq 0}} \quad \mathbb{E}_0 \int_0^\infty e^{-\rho t} \frac{c_t^{1-\sigma}}{1-\sigma}\, dt$$

subject to

$$\begin{aligned} db_t &= \Big[y^s + (1-e_t)x_t w_t + e_t \widetilde{\Pi}(z_t, b_t) + r b_t - c_t \Big] dt, \\ z_t &\in \{z_1, z_2, z_3, z_4\}, \\ x_t &\in \{0,1\}, \\ e_t &\in \{0,1\}, \\ b_t &\geq 0, \end{aligned}$$

where static entrepreneurial profits are

$$\widetilde{\Pi}(z_t, b_t) = \max_{l_t,\, 0\leq k_t \leq b_t} \left\{ z_t \big(k_t^{\alpha} l_t^{1-\alpha}\big)^{1-\nu} - w_t l_t - (r+\delta) k_t \right\}.$$

Relative to the Main Model, this specification preserves occupational choice and income risk, but removes partial irreversibility.

Aiyagari: Uncertainty + Credit Constraint

To shut down the investment channel while preserving self-insurance, we initialize households at the pre-transfer stationary distribution of the Main Model, $G_0(b_0, k_0, z_0, x_0)$, and then hold productive capital and occupational choices fixed at their inherited levels. After the transfer, the household solves

$$\max_{\{c_t\}_{t\geq 0}} \quad \mathbb{E}_0 \int_0^\infty e^{-\rho t} \frac{c_t^{1-\sigma}}{1-\sigma}\, dt$$

subject to

$$
\begin{aligned}
db_t &= \Big[y^s + (1-e_0)x_t w_t + e_0\Pi(z_t,k_0) + rb_t - c_t\Big]dt,\\
dk_t &= 0,\\
z_t &\in \{z_1,z_2,z_3,z_4\},\\
x_t &\in \{0,1\},\\
b_t &\geq 0, \qquad k_t = k_0, \qquad e_t = e_0.
\end{aligned}
$$

Thus, the transfer changes only liquid wealth. Income remains risky through wage income $x_t w_t$ and entrepreneurial profits $\Pi(z_t,k_0)$, but the household can no longer use the transfer to expand capital and raise future profits.

No Shocks: Certainty + Credit Constraint

The simplest benchmark shuts down both uncertainty and entrepreneurship. Setting $e_t = 0$, $x_t = 1$, the household solves

$$
\max_{\{c_t\}_{t\geq 0}} \quad \int_0^\infty e^{-\rho t}\frac{c_t^{1-\sigma}}{1-\sigma}\,dt
$$

subject to

$$
\begin{aligned}
db_t &= (y^s + w + rb_t - c_t)\,dt,\\
b_t &\geq 0.
\end{aligned}
$$

This benchmark shuts down the role of risk and the investment motive, leaving only pure consumption smoothing under a borrowing constraint. As shown by Park (2006), this model admits a simple closed-form characterization. Under the standard impatience condition ($\rho > r$), the household converges to the stationary allocation $c = y^s + w$ and $b = 0$, so the transfer affects only transition dynamics and not long-run outcomes.

Table 7: **Transition Dynamics: Decay Statistics by Model**

	Main Model			Reversible			Aiyagari			No Shocks		
Outcome	Impact	t_{50}	t_{90}	Impact	t_{50}	t_{90}	Impact	t_{50}	t_{90}	Impact	t_{50}	t_{90}
Wealth (% ss output)	62.8	3.5	20.3	62.8	6.7	26.0	62.8	2.5	7.7	62.8	1.8	4.2
Consumption (% ss)	16.1	3.0	11.1	9.6	6.2	23.7	17.0	2.9	8.2	20.9	3.0	5.5

Notes: This table compares impulse-response behavior across models under partial equilibrium, holding prices fixed at their stationary-equilibrium values. Outcomes are measured relative to a counterfactual economy with no transfer. *Impact* denotes the response on impact; t_{50} and t_{90} denote the years by which 50 and 90 percent of the initial impact have dissipated, respectively. Consumption is reported as percent deviation from the no-transfer consumption path, and wealth as the difference from the no-transfer path scaled by steady-state gross output.

Table 7 shows how aggregate consumption and wealth evolve after the initial transfer. Because the initial increase in wealth is the same across models by construction, the informative moments in Table 7 are the impact consumption response and the decay horizons.

Impact consumption falls from 20.9 percent in the No Shocks benchmark to 17.0 percent in Aiyagari, 16.1 percent in the Main Model, and 9.6 percent in the Reversible model. Comparing No Shocks to Aiyagari isolates the effect of precautionary saving: once income is risky, households spend less of the transfer on impact. The move from Aiyagari to the Main Model is small because irreversibility dampens the immediate investment response, so impact consumption in the Main Model is only slightly below Aiyagari. The large drop comes in the Reversible model, where households can reallocate the transfer into productive capital immediately and unwind that decision later if conditions change.

The persistence statistics show that precautionary saving matters, but that the investment channel is the dominant source of long-lived effects. The t_{90} horizon is longest in the Reversible model (23.7 years for consumption, 26.0 for wealth), where households freely convert the transfer into productive capital. Uncertainty by itself generates only a modest buffer-stock response, whereas the ability to transform the transfer into productive assets creates a far more persistent propagation mechanism.

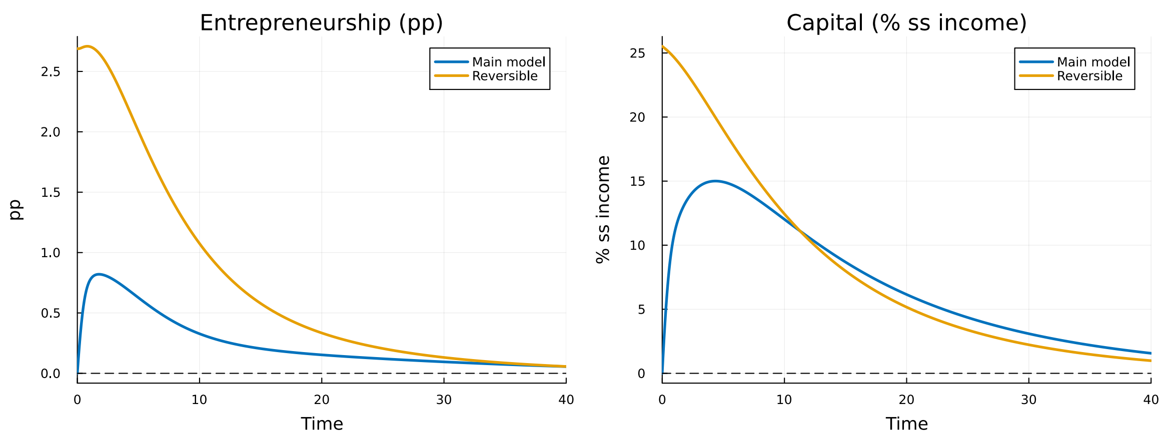

Figure 7: The figure compares the impulse responses to the cash transfer in the main model and in the model without irreversibility, under partial equilibrium with prices fixed at their pre-transfer stationary-equilibrium levels. Entrepreneurship is reported as the percentage-point deviation from steady state, while capital is reported as the deviation from steady state as a share of steady-state income.

The Main Model exhibits a distinctive pattern when compared to the Reversible case. Its half-life is short, resembling Aiyagari more than Reversible (3.0 versus 2.9 years for consumption, 3.5 versus 2.5 for wealth), which would suggest a weak investment channel early on. Yet its t_{90} is much longer than Aiyagari (11.1 versus 8.2 for consumption, 20.3 versus 7.7 for wealth) and not far below the Reversible model's long tail. In the Reversible case, a long half-life and a long tail go together; in the Main Model they come apart. Irreversibility is the reason. Figure 7 contrasts the dynamics of entrepreneurship and productive capital in the two models: in the Reversible case both respond instantly on impact, whereas in the Main Model they build up gradually along a hump-shaped path. Irreversibility makes entering entrepreneurship a slow and cautious process. Few households invest early, so the initial decay resembles the Aiyagari economy, where the investment channel is absent. But those who do enter accumulate slowly and cannot easily reverse course; their capital lingers long after the transfer, producing a tail in aggregate wealth and consumption much closer to the Reversible model than the short half-life would suggest. Irreversibility not only slows entry but also slows exit.

8. Domestically Financed Intervention

TBA

9. Conclusion

TBA

References

Asker, J., Collard-Wexler, A., De Loecker, J., 2014. Dynamic inputs and resource (mis)allocation. Journal of Political Economy 122, 1013–1063. doi:10.1086/677072.

Balboni, C., Bandiera, O., Burgess, R., Ghatak, M., Heil, A., 2022. Why do people stay poor? The Quarterly Journal of Economics 137, 785–844. doi:10.1093/qje/qjab045.

Baley, I., Blanco, A., 2026. The macroeconomics of irreversibility. The Review of Economic Studies , rdag001doi:10.1093/restud/rdag001.

Bandiera, O., Burgess, R., Das, N., Gulesci, S., Rasul, I., Sulaiman, M., 2017. Labor markets and poverty in village economies. The Quarterly Journal of Economics 132, 811–870. doi:10.1093/qje/qjx003.

Banerjee, A., Duflo, E., Goldberg, N., Karlan, D., Osei, R., Pariente, W., Shapiro, J., Thuysbaert, B., Udry, C., 2015. A multifaceted program causes lasting progress for the very poor: Evidence from six countries. Science 348, 1260799. doi:10.1126/science.1260799.

Banerjee, A., Duflo, E., Sharma, G., 2021. Long-term effects of the targeting the ultra poor program. American Economic Review: Insights 3, 471–486. doi:10.1257/aeri.20200667.

Banerjee, A., Karlan, D., Osei, R., Trachtman, H., Udry, C., 2022. Unpacking a multi-faceted program to build sustainable income for the very poor. Journal of Development Economics 155, 102781. doi:10.1016/j.jdeveco.2021.102781.

Barker, N., Karlan, D., Udry, C., Wright, K., 2023. The Fading Treatment Effects of a Multi-Faceted Asset-Transfer Program in Ethiopia. Working Paper. Buffett Institute Global Poverty Research Lab. URL: https://ssrn.com/abstract=4514167, doi:10.2139/ssrn.4514167.

Bauchet, J., Morduch, J., Ravi, S., 2015. Failure vs. displacement: Why an innovative anti-poverty program showed no net impact in South India. Journal of Development Economics 116, 1–16. doi:10.1016/j.jdeveco.2015.03.005.

Bryan, G., Chowdhury, S., Mobarak, A.M., 2014. Underinvestment in a profitable technology: The case of seasonal migration in bangladesh. Econometrica 82, 1671–1748. doi:10.3982/ECTA10489.

Buera, F.J., Kaboski, J.P., Shin, Y., 2011. Finance and development: A tale of two sectors. American Economic Review 101, 1964–2002. doi:10.1257/aer.101.5.1964.

Buera, F.J., Kaboski, J.P., Shin, Y., 2014. Macro-perspective on asset grants programs: Occupational and wealth mobility. American Economic Review: Papers & Proceedings 104, 159–164. doi:10.1257/aer.104.5.159.

Buera, F.J., Kaboski, J.P., Shin, Y., 2015. Entrepreneurship and financial frictions: A macrodevelopment perspective. Annual Review of Economics 7, 409–436. doi:10.1146/annurev-economics-080614-115348.

Buera, F.J., Kaboski, J.P., Shin, Y., 2021. The macroeconomics of microfinance. The Review of Economic Studies 88, 126–161. doi:10.1093/restud/rdaa047.

Buera, F.J., Kaboski, J.P., Townsend, R.M., 2023. From micro to macro development. Journal of Economic Literature 61, 471–503. doi:10.1257/jel.20211537.

Buera, F.J., Shin, Y., 2013. Financial frictions and the persistence of history: A quantitative exploration. Journal of Political Economy 121, 221–272. doi:10.1086/670271.

Dupas, P., Robinson, J., 2013. Savings constraints and microenterprise development: Evidence from a field experiment in kenya. American Economic Journal: Applied Economics 5, 163–192. doi:10.1257/app.5.1.163.

Egger, D., Haushofer, J., Miguel, E., Niehaus, P., Walker, M.W., 2022. General equilibrium effects of cash transfers: Experimental evidence from kenya. Econometrica 90, 2603–2643. doi:10.3982/ECTA17945.

Egger, D., Walker, J., Mughogho, W., Graff, T., 2024. General equilibrium effects of universal cash transfers: Experimental evidence at-scale from malawi. URL: https://www.dennisegger.net/. work in progress. Accessed: 2026-04-17.

Fujimoto, J., Lagakos, D., VanVuren, M., 2023. Macroeconomic Effects of 'Free' Secondary Schooling in the Developing World. NBER Working Paper 31029. National Bureau of Economic Research. doi:10.3386/w31029. previously circulated as "Aggregate and Distributional Effects of 'Free' Secondary Schooling in the Developing World".

Giné, X., Townsend, R.M., 2004. Evaluation of financial liberalization: A general equilibrium model with constrained occupation choice. Journal of Development Economics 74, 269–307. doi:10.1016/j.jdeveco.2003.03.005.

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

Jeong, H., Townsend, R.M., 2007. Sources of tfp growth: Occupational choice and financial deepening. Economic Theory 32, 179–221. doi:10.1007/s00199-006-0201-8.

Kermani, A., Ma, Y., 2023. Asset specificity of nonfinancial firms. The Quarterly Journal of Economics 138, 205–264. doi:10.1093/qje/qjac030.

de Mel, S., McKenzie, D., Woodruff, C., 2009. Measuring microenterprise profits: Must we ask how the sausage is made? Journal of Development Economics 88, 19–31. doi:10.1016/j.jdeveco.2008.01.007.

Midrigan, V., Xu, D.Y., 2014. Finance and misallocation: Evidence from plant-level data. American Economic Review 104, 422–458. doi:10.1257/aer.104.2.422.

Moll, B., 2014. Productivity losses from financial frictions: Can self-financing undo capital misallocation? American Economic Review 104, 3186–3221. doi:10.1257/aer.104.10.3186.

Park, M.H., 2006. An analytical solution to the inverse consumption function with liquidity constraints. Economics Letters 92, 389–394. doi:10.1016/j.econlet.2006.03.019.