Slackness, Openness, and the Anatomy of Cash Transfer Multipliers*

Minki Kim†
University of Mannheim

Mitchell VanVuren‡
Vanderbilt University

April 2026

Preliminary draft, for conference submission

Abstract

This paper develops a general equilibrium macroeconomic model to rationalize why large-scale cash transfers in low-income settings generate high multipliers with little price inflation. We provide two mechanisms, slackness in local production and openness to external trade, and show they have distinct policy implications. We calibrate the model to baseline data from a large-scale cash-transfer experiment in Malawi. The calibrated economy lies on a slackness plateau where firms have substantial idle capacity, muting the price response and generating a large local GDP multiplier. Within this regime, the welfare-maximizing transfer design spreads transfers across more villages rather than concentrating them at higher per-household amounts. We also explore additional mechanisms that may shape transfer multipliers, with predictions the Malawi endline can test directly.

*Financial support from the Structural Transformation and Economic Growth (STEG) is gratefully acknowledged. This paper is a part of the STEG's initiative, "The Macroeconomic Impacts of Cash Transfers."
†Email: minki.kim@uni-mannheim.de
‡Email: mitchell.vanvuren@vanderbilt.edu

1 Introduction

Cash transfers have become one of the most widely used social safety net programs in the developing world. Understanding the general equilibrium effects of these programs is especially important in low-income settings, where large-scale transfers impose significant fiscal burdens. A growing body of empirical work has examined this question through large-scale randomized controlled trials. What is surprising is that previous RCT studies have found little evidence of price effects while finding positive local transfer multiplier effects, where total consumption and income increase by more than the value of the transfer (Cunha, De Giorgi and Jayachandran, 2019; Egger, Haushofer, Miguel, Niehaus and Walker, 2022). Yet the structural mechanisms behind these patterns, and what they imply for transfer design, remain unsettled.

What drives these small price effects? One possibility is that there is "slackness" in the utilization of production inputs (Lewis, 1954), i.e., the labor markets in developing countries may feature perpetually dormant productive capacity due to market failure (Walker, Shah, Miguel, Egger, Soliman and Graff, 2024). In this case, transfers can activate this dormant capacity, allowing supply to expand with little increase in price. A second explanation, which we call "openness", stems from the macroeconomic literature that highlights the importance of trade in determining local responses to shocks (Farhi and Werning, 2016). In a highly open village that is connected to and trades extensively with external markets, the increase in demand from transfers can be met through increased consumption of goods from outside the village. As a result, such a village would experience little increase in local prices as external markets are able to act as an additional source of supply to meet increasing demand.

The goal of this paper is to distinguish between these two competing explanations for the limited price effects of cash transfers. This question is not only intellectually interesting but also important for guiding transfer implementation, as the two mechanisms have distinct implications for how cash transfers should be targeted and the potential risks in scale-up. If slackness dampens price impacts and drives large multiplier effects, transfers should be directed toward more isolated, demand-constrained areas where slack is likely to exist. However, slackness also implies limits to scale-up: sufficiently large transfers could exhaust slackness and trigger price increases. In this case, scaling by providing smaller transfers that fall below the slackness threshold to more villages will lead to larger aggregate welfare gains.

On the other hand, the policy implication differs significantly if openness drives the results. In this case, transfers should target less isolated areas, rather than more isolated areas, as villages with greater trade access can draw on a larger pool of external supply to meet rising local demand. An openness-dominated world also presents different risks to scale-up. Rather than the risk that transfers are too large and overwhelm residual slackness, the concern is that distributing transfers too broadly may eventually exhaust external supply, triggering price increases with limited welfare gains.

In this paper, we develop a general equilibrium macroeconomic model that incorporates both slackness and openness, which can be estimated using the results of an experimental evaluation of cash transfers that is currently being conducted. The model has several features relevant to studying the quantitative importance of slackness and openness. First, the model features a fully flexible price level in order to accommodate the impact of cash transfers on prices and the (extreme) potential outcome that increases in money simply translate to increases in prices without any increase in production or consumption. Second, there is inter-village trade, which enables external supply to absorb increased demand. Additionally, the model incorporates slackness through what we call "demand-constrained production," where firms produce only when they have customers.[1]

Although slackness and openness imply different policy prescriptions, both likely operate to some degree. The key question is their relative quantitative importance. The model provides insight that the size of the *income* multiplier of transfers is particularly informative to distinguish the two explanations. The intuition is straightforward—it comes from the fact that with trade, consumption and income multipliers might not be identical. In a world where slackness dominates, local firms can simply activate the dormant production capacity to cater to higher demand triggered by the transfer. This would lead to higher income for the local households, leading to a large income multiplier. In contrast, when openness dominates, the increased demand is met primarily through foreign supply, leaving local incomes relatively unchanged. Thus, while both scenarios may generate high *consumption* multipliers, slackness would yield much higher *income* multipliers.

We bring this model to baseline data from the GiveDirectly cash-transfer experiment in Malawi, where transfers of approximately \$550 per adult (100–150% of annual local GDP in most treated villages) are delivered with timing randomized at both household and village levels. We calibrate five household and firm parameters to four Malawi moments and the marginal propensity to consume (MPC) from Egger et al. (2022). The calibrated econ-

[1] Imagine, for example, a grain miller. In a low-demand environment, they may only have a few customers each day, and processing this grain may only take an hour or two of labor. However, because they do not know when these customers will arrive, they must tend their shop the whole day. Should their number of customers double, the miller could easily double their production. In this sense, the miller's production is completely demand-constrained.

omy indicates that the experimental sample lies on what we call the *slackness plateau*, a regime in which non-agricultural firms operate with substantial idle capacity and the price response is fully muted.

The model predicts a high GDP multiplier of 1.83 for the Malawi experiment, although lower than the 2.58 reported by Egger et al. (2022) for the comparable Kenyan experiment. The welfare-maximizing transfer design spreads transfers across more villages rather than concentrating them at higher per-household amounts, because pushing per-household intensity beyond the slackness plateau threshold delivers diminishing welfare gains as local prices begin to rise. We also propose a candidate explanation for the gap between the model's prediction and the Kenya estimate. Transfer-induced demand draws labor from agriculture into non-agricultural production, but a competition externality among non-agricultural producers, who compete for the same pool of buyers by staying open longer, depresses the aggregate marginal product of that labor. The implied reallocation is too small to be identified in the existing experimental data, but the Malawi endline may be used to measure cross-sector labor reallocation directly and to test the model's prediction.

Related Literature This paper highlights the opportunity for productive collaboration between experimental results and a structural, general equilibrium macro model. Our model allows us to leverage the experimental moments to generate quantitative comparisons between the strength of channels and benchmark the aggregate impacts of various potential policies. In doing so, we contribute to a growing literature in macroeconomic development that integrates experimental data with structural modeling to study the general equilibrium effects of development policies (Todd and Wolpin, 2006; Akcigit, Alp and Peters, 2021; Buera, Kaboski and Townsend, 2023; Fujimoto, Lagakos and VanVuren, 2023; Lagakos, Mobarak and Waugh, 2023).

We provide an explanation for the large fiscal multipliers in low-income settings, which has been documented in several field experiments (Cunha et al., 2019; Egger et al., 2022; Banerjee, Faye, Krueger, Niehaus and Suri, 2023; Mendes, Miyamoto, Nguyen, Pennings and Feler, 2023). While slackness in production has been proposed as one explanation (Walker et al., 2024), we introduce openness as an alternative channel and develop a model that enables a quantitative comparison between these mechanisms.

This paper is also related to the macroeconomic literature estimating the local and aggregate fiscal multipliers using regional variations (Nakamura and Steinsson, 2014; Ramey and Zubairy, 2018; Hazell, Herreno, Nakamura and Steinsson, 2022). Since cash transfers in our setting are purely external injections without an offsetting reduction in spending in other regions, they offer a clean setting for estimating multipliers. We use the observed multipliers as a tool to isolate the underlying mechanism behind the price effects of fiscal stimulus.

2 Model

There are an arbitrary number of villages N, each populated by a unit measure of households. Households split their time between working in agriculture and producing their unique variety of non-agricultural good. Each household consumes an aggregate consumption good created by combining agricultural goods and non-agricultural varieties from all villages.

Preferences: Time is discrete. A household in village h (for "home") exhibits typical CRRA preferences with respect to aggregate consumption C.

$$u(\{C_t\}_{t=0}^{\infty}) = \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\rho}}{1-\rho} \tag{1}$$

Aggregate consumption C is created by combining agricultural goods C_A and household-village-specific varieties of non-agricultural goods $C_{N,v}(i)$ (i is a household index) via three constant elasticity of substitution layers.

First, non-agricultural household-village varieties are aggregated into village varieties via a standard CES aggregator.

$$C_{N,v} = \Big(\int C_{N,v}(i)^{\frac{\sigma-1}{\sigma}} di \Big)^{\frac{\sigma}{\sigma-1}} \tag{2}$$

Second, village varieties are aggregated with a constant elasticity of substitution σ_f.

$$C_N = \left(\sum_{i=1}^{N} C_{N,i}^{\frac{\sigma_f - 1}{\sigma_f}} \right)^{\frac{\sigma_f}{\sigma_f - 1}} \tag{3}$$

Finally, the non-agricultural aggregate and agricultural goods are combined.

$$C = [\alpha^{1/\sigma_a} C_A^{\frac{\sigma_a - 1}{\sigma_a}} + (1-\alpha)^{1/\sigma_a} C_N^{\frac{\sigma_a - 1}{\sigma_a}}]^{\frac{\sigma_a}{\sigma_a - 1}} \tag{4}$$

Trade Costs: Agricultural goods are freely traded at a fixed price (due to international markets) and serve as the numeraire. We denote the price of non-agricultural goods of household i in village v as $p_{N,v}(i)$. Villages face iceberg trade costs in their interactions with each other. We denote the trade cost that village v faces when purchasing from village w as $\tau_{v,w}$ so that (for example) village v's effective price of the household i in village w's good is $\tau_{v,w}p_{N,w}(i)$.

Saving and Budget: Households have access to a risk-free asset a that they can use to save but not borrow ($a \geq 0$) that pays exogenous return R. We allow for the possibility that $R < 1$ so that this represents a (net) costly saving technology. Their budget depends on their consumption as well as their income I (discussed below) and a potential cash transfer T.

$$a' + \sum_{v=1}^{N} \int_{i} {}_{h,v} p_{N,v}(i) C_{N,v}(i)\, di + C_A = Ra + I + T \tag{5}$$

where h is the home village of the consumer.

In order to match data on the aggregate marginal propensity to consume, we allow for a share θ of households to be entirely hand-to-mouth, choosing $a' = 0$ each period, rather than optimally smoothing consumption. Without this parameter, the model delivers MPCs much lower than those measured in cash transfer experiments and, consequently, severely underpredicts multipliers.

Non-Agricultural Production: Our model of non-agricultural production draws from Walker et al. (2024) who emphasize the possibility that production is "slack" in developing countries. We embed this concept in our model via a theory of "demand-constrained" production.

Consider, as an example, a retailer operating in one of these villages. Despite attending their shop or stall the entire day, this retailer may only have a handful of customers and spend the day otherwise idle. Should their number of customers double, the retailer could double production with no additional labor input, though they would require more intermediate inputs. In this sense, the retailer's production is entirely constrained by demand (at least at the margin) and supply is slack.

At the same time, despite not needing all of their labor input to serve customers, the retailer cannot costlessly allocate their time to some other use (e.g. agricultural production in the context of the model). The exact arrival times of customers are unpredictable and, if the retailer is not open when they arrive, they will be unable to serve them. Thus if the retailer reduces their labor input by half and closes their business half the time, they can also expect half the customers.[2] As a consequence, even if supply is slack, as above, the retailer may not allocate time away from their business, and this slackness may persist even in equilibrium.

To implement these ideas in our model, we begin with a simple linear production function. Each household is born with some fixed productivity z drawn from some (bounded) distribution $G(z)$. The household's production technology uses a Cobb-Douglas combination of agricultural goods and the local non-agricultural aggregate as intermediate inputs with share parameter ω and converts $\frac{1}{z}$ units of this combined input into a unit of the household's unique variety. For simplicity, we define $c(z) = \frac{1}{z}$ so that $1^{1-\omega}P_{N,v}^{\omega}c(z)$ is the cost-minimizing per-unit cost of production for a z-type household in village v ($P_{N,v}^{\omega}$ is the CES ideal price index for the village's aggregate non-agricultural good). Profits at a given quantity-price (q, p) combination are then given by

$$\pi(p, q) = pq - P_{N,v}^{\omega}c(z)q \tag{6}$$

Although the household would like to produce as much output as possible, it faces two restrictions. First, each household is only able to serve $\bar{q}$ customers per unit of time so that quantity cannot exceed $\bar{q}l_N$ where l_N is the amount of time allocated to non-agricultural production.

$$q_{N,v}(i) \leq \bar{q}l_N(i) \tag{7}$$

Second, the household is constrained by demand for its good and cannot sell more than the quantity demanded at a given price. Within each period, time flows continuously and each buyer visits the local non-agricultural market exactly once. In reality, not every consumer visits every village market, but we can think of trade between villages as occurring through traders who inherit their destination market's preferences and visit each village once per period.

The arrival times of buyers are idiosyncratic, unpredictable (from the perspective of the household engaged in production), and uniformly distributed in time across the period.[3] Upon arrival, buyers can only purchase from firms that are currently open so that a firm that is open for a fraction $l_N(i)$ of the period is therefore found open by a fraction $l_N(i)$ of buyers.

[2] In reality, demand is somewhat predictable and the ability to close during the least productive hours suggests that this elasticity could be less than unity. Still, we will maintain the assumption of unit elasticity for simplicity

[3] For local consumers, the timing of market visits is shaped by their own agricultural and domestic tasks. For traders, it is shaped by travel logistics.

Conditional on finding firm i open, a buyer allocates spending across available varieties according to their CES demand described in (2). We assume that a household who spends share $l_N(i)$ of their time in non-agricultural production has their firm open for share $l_N(i)$ of the period, with the exact opening hours distributed randomly and uniformly throughout the period. Thus the share of goods/firms available to a buyer is independent of their arrival time. Each buyer therefore spends share $p_{N,v}(i)^{1-\sigma} / \int l_{N,v}(j)\, p_{N,v}(j)^{1-\sigma} dj$ of their total village-v non-agriculture spending on firm i (conditional on i being open).

Since a fraction $l_{N,v}(i)$ of buyers find firm i open, aggregating across all buyers gives

$$q_{N,v}(i) = l_{N,v}(i)\, p_{N,v}(i)^{-\sigma} \frac{S_{N,v}}{\int l_{N,v}(j) p_{N,v}(j)^{1-\sigma} dj} \tag{8}$$

$$S_{N,v} = \sum_{w=1}^{N} \int_i s_{N,w,v} I_w(i) di \tag{9}$$

where $s_{N,v,w}$ is the share of income spent by households in village w on non-agricultural goods from village v so that $S_{N,v}$ is total spending on v's non-agricultural goods.[4]

This demand curve captures both key features of slackness described above. If the household reduces its labor input (hours open) by half, demand for its goods falls by half because fewer buyers find the firm open. Equation (7) captures the other key feature: if the optimal quantity $q^*_{N,v}(i)$ is below $\bar{q} l^*_N(i)$, the household can accommodate increases in demand (i.e. increases in $S_{N,v}$) without additional labor input.

Agricultural Production: Agricultural production is comparatively simple. Each household operates a decreasing returns-to-scale agricultural production technology with common productivity A that takes only labor as an input. Agricultural production is then given by

$$q_{A,v} = \int_i A l_{A,v}(i)^{\gamma}\, di \tag{10}$$

and is sold at a price of unity under perfect competition.

Time Constraint and Income Maximization: Each household is endowed with one unit of time and receives no utility from leisure so that their time constraint is given by

$$l_{N,v}(i) + l_{A,v}(i) = 1 \tag{11}$$

[4] Expressions for these shares, in terms of parameters and prices, follow directly from the household's consumption optimization problem but are tedious and relegated to the appendix.

The income maximization problem of the household is then given by

$$\max_{q_N, p_N, l_N, l_A} \pi(p_N, q_N) + A l_A^{\gamma} \tag{12}$$

subject to (7) - (9) and (11), where the village and household indices v and i have been suppressed for readability.

Typically in CES demand systems, the optimal price takes the form of a constant markup over marginal cost. This is also the case here, under the condition that equation (7), which dictates the maximum production of a household, does not bind. In the case where this constraint is binding, the household simply increases its price (even beyond this markup) in order to bring demand back down to its maximum production.

$$p^* = \begin{cases} \frac{\sigma}{\sigma-1} P^{\omega}_{N,v} c(z) & \text{if} \quad \left(\frac{\sigma}{\sigma-1} P^{\omega}_{N,v} c(z)\right)^{-\sigma} \frac{S_{N,v}}{\int l_{N,v}(j) p_{N,v}(j)^{1-\sigma} dj} \leq \bar{q} \\ \left(\bar{q}^{-1} \frac{S_{N,v}}{\int l_{N,v}(j) p_{N,v}(j)^{1-\sigma} dj}\right)^{\frac{1}{\sigma}} & \text{otherwise} \end{cases} \tag{13}$$

The household also equates the marginal product of agricultural and non-agricultural labor (the latter of which depends on the chosen price p^*).

$$A\gamma {l^*_A}^{\gamma-1} = \left(p^* - P^{\omega}_{N,v} c(z)\right)\left(p^{*-\sigma} \frac{S_{N,v}}{\int l_{N,v}(j) p_{N,v}(j)^{1-\sigma} dj}\right) \tag{14}$$

Idiosyncratic Shocks and Income: Finally, we allow households to experience idiosyncratic and symmetric shocks to both agriculture and non-agricultural productivity so that a household's total income with shock y is given by

$$I(y) = \left(\pi(p_N, q_N) + A {l_A}^{\gamma}\right) y \tag{15}$$

For simplicity, we specialize to the binary Markov case with transition matrix M and normalize the high-productivity state to a value of $y_h = 1$ so that the low state reflects a temporary, proportional productivity decline of y_l.

Equilibrium: Though we relegate the formal definition of recursive competitive equilibrium to the appendix, it is worth describing the market clearing conditions. The first set of conditions corresponds to the markets for the households' unique non-agricultural varieties. The demand-curve constraint in equation (8) ensures the household will always choose a pair (p_N, q_N) that lies on their individual demand curve, conditional on perceiving aggregate demand correctly. Thus market clearing for non-agriculture simply requires that perceptions of aggregate income align with reality.

$$I_v(i) = \big(\pi(p_N^*, q_N^*) + A{l_A^*}^{\gamma}\big)y(i) \tag{16}$$

Again, it is worth noting that while we have suppressed it for simplicity, the values on the right-hand side of this equation can differ across villages and households — a fact we emphasize by indexing the left-hand side by v and i.

Together, the household consumption optimization problem (described by 1-5), the household income maximization problem (described by 6-14), and the market clearing condition (described by 16) fully determine all endogenous variables in competitive equilibrium.

3 Qualitative Examples

How do slackness and openness affect the impact of transfers on local prices, output, and consumption? Here we build some intuition using an example economy with $N = 2$ villages with one village corresponding to the treated village receiving cash transfers and the other corresponding to control. For simplicity, we also assume that the control village is *large* relative to the treated village, so the outcomes in the treated village have no impact on prices in the control village.

We start from a baseline economy featuring no slackness and little amount of trade (due to very high trade costs) and consider two changes. The first (the "Slack Economy") maintains the small amount of trade but introduces a positive amount of slackness. The second (the "Open Economy") maintains the zero-slack assumption but reduces trade costs to generate a substantial amount of trade.[5] Figure 1 compares the impact of cash transfers on the treated village between the Slack Economy and the Baseline (column a) and the Open Economy and the Baseline (column b). We vary the size of the transfer from 0 to 30 percent of average village income and plot the impact on local non-agriculture prices as well as local real GDP, providing insight into how the impact of the transfer varies with its size.

Beginning with column (a), which compares the impact of the transfer between the Slack economy and the Baseline, we see that the presence of slackness substantially mitigates the price impacts of the transfer. Though agricultural prices (not shown) do start adjusting

[5] It is important to highlight that these results come from a simplified economy with parameters chosen essentially at random. A full calibrated model is introduced in Section 4.

Figure 1: Impact of Transfers on Slack vs Open Economies

(a) Slack Economy vs Baseline | (b) Open Economy vs Baseline

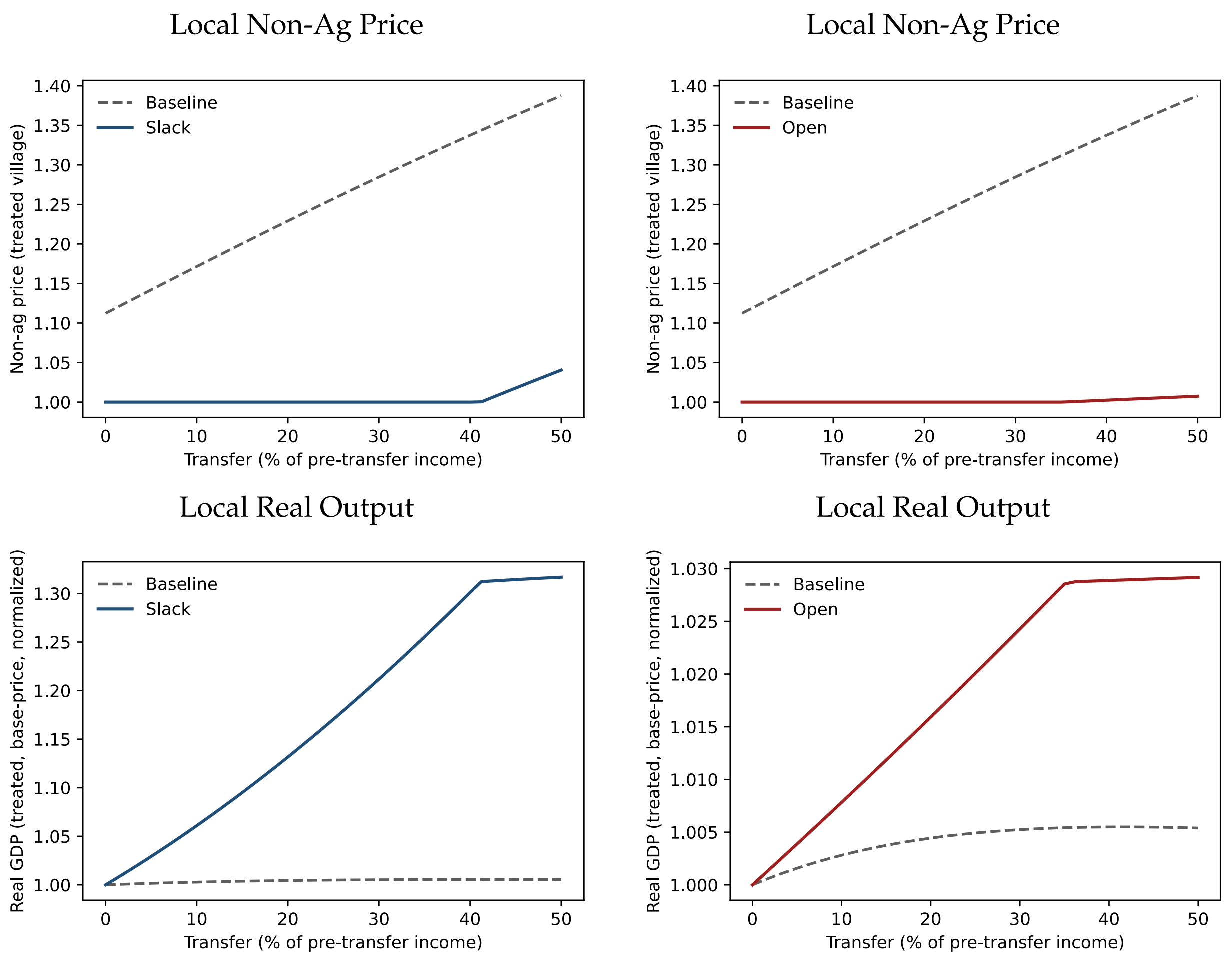

immediately, non-agricultural prices stay completely flat until the transfer is large enough to overcome underlying slackness, at which point they also begin to appreciate. As a result, local output increases dramatically up until the point where all slackness has been eliminated. At this point, the GDP gains from further increasing the size of the transfer fall off dramatically, though they are still positive as a result of the fact that even these economies feature some non-zero amount of trade.[6]

Column (b) displays the same outcomes for the Open economy, compared to the Baseline. As with slackness, openness mitigates the increase in local prices that occur as a result of the transfer. Here, this occurs because competition with foreign producers lowers the extent to which locals can increase prices. As a result, the real value of the transfer represents

[6] A technical necessity to avoid an indeterminate local price level.

Figure 2: Consumption and Income Multipliers in Slack vs Open Economies

(a) Slack Economy vs Baseline (b) Open Economy vs Baseline

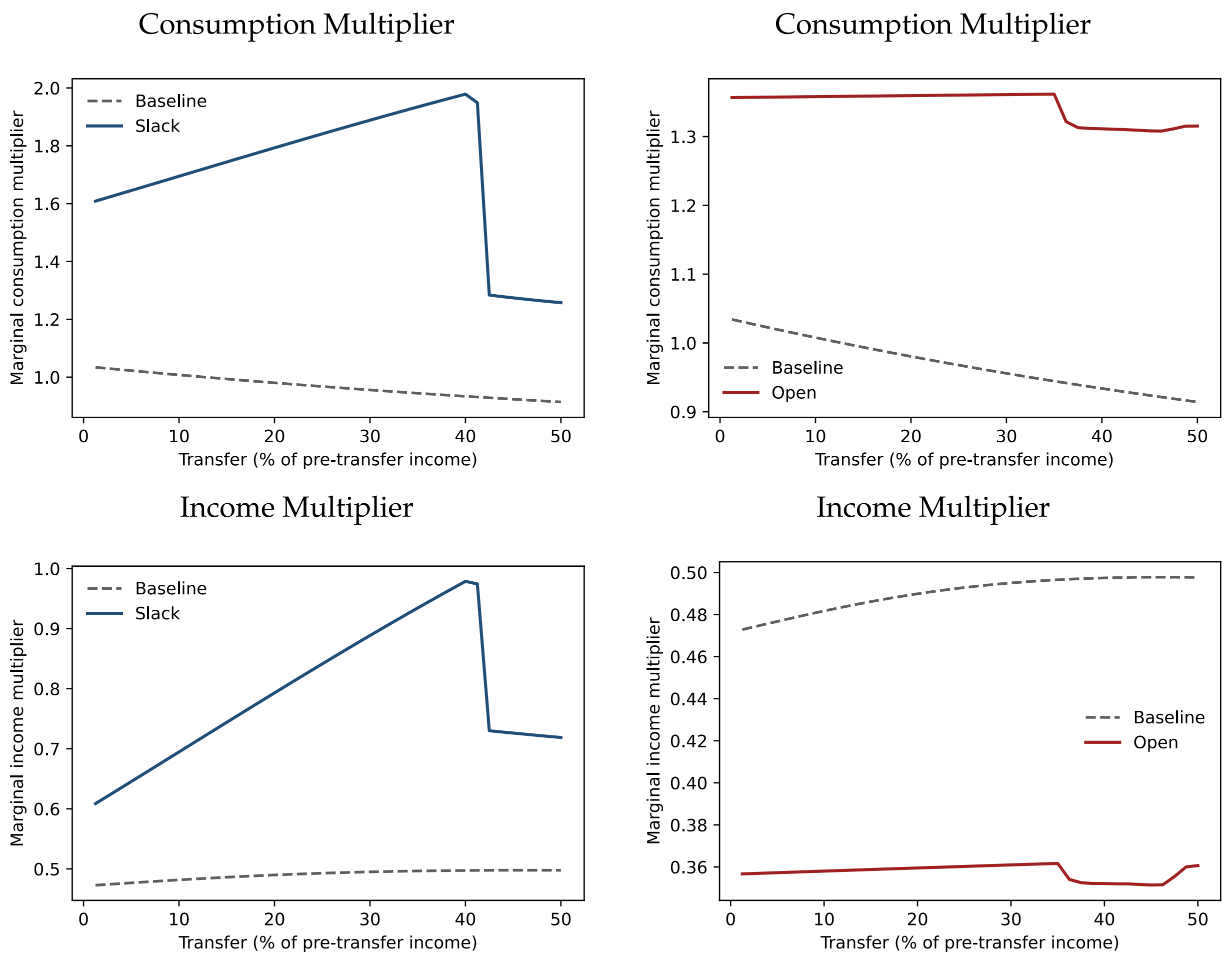

a larger local demand shock and leads to a large increase in local output.

Differences between Slackness and Openness: The results in Figure 1 suggest that slackness and openness operate similarly and can both explain the muted price effects observed in past experiments. It is tempting to use this to conclude that distinguishing between the two is irrelevant other than perhaps as a curiosity. This, however, turns out not to be the case. The two have very different properties that directly impact how potential transfers should be structured to maximize their impact on welfare.

The distinction between the two is seen most clearly by comparing their impacts on local multipliers. We focus in particular on the distinction between consumption multipliers (the multiple by which local consumption increases in response to a transfer) and income multipliers (the multiple by which local income increases). These are not necessarily identical because each village produces a unique variety with differences in local and foreign prices. Income depends entirely on local prices while consumption depends on both local and foreign prices. In essence, there is something akin to a "terms-of-trade" effect that separates the two.

Figure 2 displays how slackness (column a) and openness (column b) impact the local consumption (first row) and income (second row) multipliers. To emphasize how these multipliers vary with the size of the transfer, we plot the marginal (rather than average) multiplier so that, for example, the value displayed at 30 percent represents the additional consumption/income that would be gained from increasing the size of the transfer from 29 percent of GDP to 30 percent.

Compared to the baseline economy, both slackness and openness result in a large consumption multiplier. That is, a larger increase in real consumption in response to any given transfer size. The two do, however, make different predictions for how the multiplier varies with the size of the transfer. In the slack economy, smaller transfers have much larger multipliers, but as the size of the transfer increases and eventually overwhelms the underlying slackness, the multiplier falls closer to the baseline economy. The open economy, on the other hand, maintains a higher, positive multiplier even in the presence of large transfers, highlighting the difference in what these two channels imply about the extent to which transfers can continue to improve welfare as their size is scaled.

Unlike the consumption multiplier, the two channels impact the income multiplier in completely opposite ways. That slackness increases the income multiplier is unsurprising at this point and, for the same reasons as the consumption multiplier above, the increase falls off dramatically as the transfers grow larger. What merits more explanation is that openness *decreases* the income multiplier. The intuition for this is very classical; a higher degree of openness means a lower marginal propensity to consume *locally* out of income, decreasing the local multiplier for traditional Keynesian reasons. Though it is clear from column (a), it is worth explicitly noting that the open economy's larger consumption share of foreign goods nevertheless leads the transfer to result in larger welfare/consumption gains than in the baseline economy. This does, however, highlight another difference in what these two channels imply for scale-up. Under openness, expanding the scope of the transfer to include the foreign village would eliminate this advantage and reduce the consumption gains.

Identification: The results of Figure 2 also demonstrate how the degrees of slackness and openness can be pinned down in data. While both can explain flat aggregate supply curves and large consumption gains in response to transfers, the difference between (local) consumption and income multipliers distinguishes the two, with smaller differences pointing towards a larger role for slackness and larger differences pointing towards a larger role for openness. In the following section, we develop a quantitative model calibrated to the Malawi GiveDirectly RCT to diagnose which explanation is quantitatively dominant in the experimental sample.

4 Calibration and Quantitative Results

The goal of this section is to parameterize the model so that its baseline steady state matches the moments observed in the Malawi GiveDirectly cash-transfer experiment, and then to report what the calibrated model predicts about prices, output, and the income–consumption multiplier gap under the treatment. Section 4.1 describes the calibration procedure and reports the model fit. Section 4.2 reports the calibrated model's baseline predictions for the experimental treatment. Section 4.3 interprets the Malawi calibration in terms of the degree of slackness and openness, and explains what the model's predicted multiplier means in light of Egger et al. (2022)'s (larger) point estimate.

4.1 Calibration Procedure

The model's parameters fall into three sets. The first set is externally calibrated to standard literature values, including discount factors, elasticities, and productivity normalizations. The second set is computed directly from Malawi data, including the agricultural expenditure share, the local input share, the iceberg trade cost, and the parameters of the income process. The third set, governing household saving and firm pricing and capacity, is jointly estimated by SMM to match five aggregate moments. Four moments are from the Malawi data; one (the 27-month cumulative MPC) is borrowed from Egger et al. (2022)'s estimates of the Kenya GiveDirectly experiment, since the Malawi experiment is still ongoing at the time of writing.

Although all moments are jointly determined in equilibrium, we pair each parameter with its most closely associated moment to build intuition. The five SMM-estimated parameters and associated moments are: the gross interest rate R (asset-to-income ratio); the within-village variety elasticity σ (Lerner markup); the unit production cost c (non-agricultural expenditure share); the hand-to-mouth share θ (MPC, taken from Egger et al., 2022); and the non-agricultural firm capacity $\bar{q}$ (capacity-utilization rate, Malawi enterprise census).[7] All five moments are steady-state objects of the model, and the SMM is

[7] Section 4.3 explains why we identify $\bar{q}$ using the capacity-utilization rate rather than the GDP multiplier.

Table 1: Externally Chosen Parameters

Parameter	Description	Value	Source/Target
Normalizations			
z	Non-ag labor productivity	1.0	Normalization
A	Agricultural TFP	1.0	Normalization
Preferences and technology			
γ_A	Ag returns to scale (DRS)	0.55	Chen et al. (2023)
CRRA	Coefficient of relative risk aversion	1.1	Standard value
β	Discount factor	0.95	Standard value
σ_{an}	Ag/non-ag CES elasticity	2.0	Prior literature
σ_f	Across-village CES elasticity	4.0	Price survey gravity regression
Data-matched structural parameters			
α	Ag expenditure share	0.258	Baseline household survey
ω_{nonag}	First-round local VA share of non-ag revenue	0.666	Enterprise census
Var$(\log y)$	Transitory income variance	0.5	Literature
p_{hl}	Markov transition prob. (high $\to$ low)	0.25	
p_{lh}	Markov transition prob. (low $\to$ high)	0.25	
τ	Iceberg trade cost	2.15	Home share $= 0.666$
Experimental design			
Transfer amount	Per-HH transfer / consumption	0.70	Chiradzulu PAP
Saturation	Within-village HH receipt rate	1.0	Chiradzulu PAP
N	Number of villages	6	RCT design
n_{treated}	Number of treated villages	4	High-saturation arm

solved at the steady-state distribution. The model's predicted GDP multiplier, a transition-path object that does not enter the SMM objective, is computed post-estimation.

Table 1 reports the externally-calibrated and data-computed parameters. The top of the table normalizes agricultural and non-agricultural productivity (A and z). The next block lists preference and technology parameters taken from the existing literature; the discount factor β, CRRA, the agricultural returns to scale γ_A (Chen, Restuccia and Santaeulàlia-Llopis, 2023), the across-village elasticity σ_f, and the agricultural/non-agricultural elasticity σ_{an}. The third block reports parameters computed directly from Malawi data. Those moments are the agricultural expenditure share α, the local input share ω, the iceberg trade cost τ, and the parameters of the binary Markov income process. The bottom block describes the experimental spatial design ($N = 6$ villages, four treated, per-household transfer size 0.70, within-village coverage of 100%), which the model is going to replicate.

Table 2 reports the five SMM-estimated parameters alongside their target moments and

Table 2: Estimated Parameters and Moment Fit.

Parameter	Value	Moment	Data	Model	Source
R	0.995	Asset/income ratio	0.882	0.882	HH baseline
σ	1.341	Lerner markup	0.746	0.746	Enterprise census
c	0.255	Non-ag expenditure share	0.742	0.742	$1-\alpha$
θ	0.878	27-month cumulative MPC	0.93	0.93	Egger et al. Table C.I
$\bar{q}$	4.82	Capacity utilization	0.291	0.291	Enterprise census

the resulting model fit.[8] All five moments are matched exactly. The estimated hand-to-mouth share is high ($\theta = 0.878$), consistent with the limited consumption-smoothing technology available in low-income rural settings. The elasticity of substitution is low ($\sigma = 1.34$), implying the markup defined as the firm's price-cost margin ($(p - MC)/p = 1/\sigma$ under CES) of roughly 0.75, which is high but plausible for retail and service firms in this context. The implied firm capacity is $\bar{q} = 4.82$, well above the slackness-plateau threshold (Section 4.3).

4.2 Model's Predictions for the Malawi Experiment

With the calibrated model in hand, We solve the model's 27-month transition path under the experimental treatment and report cumulative outcomes. The experimental design in the model matches the spatial design of the Malawi RCT (six villages, four treated, within-village coverage of one, per-household transfer size 0.70), and the calibrated parameters are those reported in Tables 1 and 2. These headline numbers serve as the calibrated model's pre-registered predictions for the Malawi endline, as they are generated from the model calibrated to the baseline data, not fit to the (yet to be observed) endline outcomes.

Table 3 summarizes the results. The model predicts a sizable GDP multiplier from the experiment. The predicted GDP multiplier is 1.83, against Egger et al. (2022)'s estimate of 2.58 in a similar RCT in Kenya. The income–consumption multiplier gap is 0.092. Changes in the non-agricultural price, the treated-village CPI, and the agricultural price are all essentially zero.

The income–consumption multiplier gap of 0.092 is quite small. Recall from Section 3 that the gap is the diagnostic that separates slackness from openness. A small gap supports

[8] The model overpredicts the agricultural income share (0.37 versus 0.23 in the data). The gap reflects the model's focus on the non-agricultural response to transfers; agricultural income composition is not a target.

Table 3: Model Predictions for the Malawi Experiment

Quantity	Prediction
Consumption mult. (nominal)	1.601
Income mult. (nominal)	1.693
Aggregate GDP multiplier	1.829
Income–consumption gap (nominal)	0.092
$\Delta P_{\text{non-ag}}$ (%)	0.000
$\Delta \text{CPI}_{\text{treated}}$ (%)	0.000
ΔP_{ag} (%)	0.000

Notes. Cumulative outcomes over the 27-month transition path under the Malawi RCT design (six villages, four treated, within-village coverage of one, per-household transfer size 0.70). Consumption and income multipliers are nominal and economy-wide.

the slackness interpretation, while a large gap would point to an openness role above what the calibrated home expenditure share implies, or to a labor-reallocation effect with a particular sign. In the next subsection, we investigate where the Malawi calibration is located in the slackness/openness regimes.

4.3 The Slackness Plateau

The calibrated model's response to the experimental treatment depends on two structural inputs. First one is the firm capacity $\bar{q}$, which governs slackness. The second ons is the home expenditure share, which governs openness. As illustrated in Section 3, three outcomes could be used as diagnostics: the GDP multiplier, the change in CPI, and the income–consumption multiplier gap. Figure 3 reports each outcome (cumulated over the 27-month transition under the experimental treatment) with varying degree of firm capacity $\bar{q}$, holding the home expenditure share at four representative values (0.20, 0.666, 0.78, 0.95).

Three regimes appear along the firm-capacity axis. In Regime I ($\bar{q}$ below approximately 1.33), steady-state production is already at full capacity; the transfer pushes prices up sharply, the GDP multiplier is well below 1, and CPI rises. In Regime II ($\bar{q}$ approximately 1.33 to 2.33), steady-state slack exists but the transfer exhausts it; CPI peaks and then declines, and the GDP multiplier rises with $\bar{q}$. Regime III is what we call the *slackness plateau*, where $\bar{q}$ exceeds approximately 2.33. In this regime, the transfer never exhausts capacity, and the three outcomes are flat in $\bar{q}$.

The home expenditure share governs openness independently of $\bar{q}$. It primarily moves CPI: more open economies (lower home expenditure share) absorb more of the transfer-induced demand through imports, which dampens the local price response. The income–consumption multiplier gap is the cleanest openness diagnostic: it is monotone in the home expenditure share and approximately flat in $\bar{q}$ within the plateau. The income–consumption gap therefore separates slackness from openness once both are present in the same calibrated economy (Section 3).

The model implies that the slackness plateau is where the baseline calibration sits on. Above the threshold, no firm reaches its capacity ceiling at any point along the transition. The transfer-induced demand is therefore fully met through quantity adjustment, and additional headroom in $\bar{q}$ is slack that the experiment never activates. The GDP multiplier saturates and becomes insensitive to $\bar{q}$, ruling out the GDP multiplier as the identifying moment for $\bar{q}$, since every value above the plateau threshold delivers the same multiplier. We therefore identify $\bar{q}$ using the capacity-utilization rate calculated from the enterprise census. The calculated value is 0.291, drawn from a sample of 3,656 enterprises in the Malawi study area. Operators are busy with customers only 29.1% of the time; the remaining 70.9% are waiting for customers, providing direct evidence of slack. The capacity-utilization moment varies smoothly in $\bar{q}$ on the plateau, and matching it to the surveyed value of 0.291 yields $\bar{q} = 4.82$ (Table 2).

5 Optimal Policy

Given the Malawi RCT's experimental design, is the experimental transfer welfare-maximizing, or would a different split across villages and per-household intensities deliver a larger welfare gain at the same cost?

We set the policy design space as a two-dimensional grid over village coverage h and per-household transfer intensity v. The grid has six coverage levels from $h = 1/6$ to full coverage $h = 1$ and eight intensities from $v = 0.20$ to $v = 1.5$ times steady-state household income, resulting in 48 (coverage, intensity) combinations in total. For each cell, we solve a separate transition path of the calibrated model under the corresponding (h, v) scenario. Consumption-equivalent welfare is calculated using the discounted utility of treated-village consumption over a 50-period transition, net of the no-transfer baseline, and total cost is the sum of transfers over villages and time. As a benchmark, the Malawi RCT's actual allocation is $(h, v) = (2/3, 0.70)$.

Figure 4 plots cost-effectiveness across the policy grid. The horizontal axis is village cover-

Figure 3: GDP Multiplier, CPI Change, and Multiplier Gap with Varying Firm Capacity $\bar{q}$.

(a) GDP Multiplier

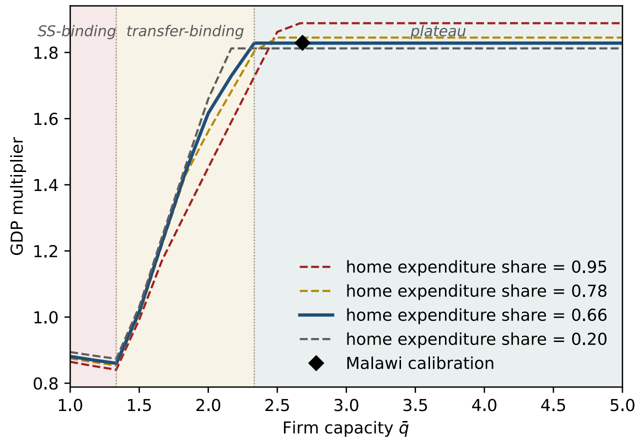

(b) Change in CPI

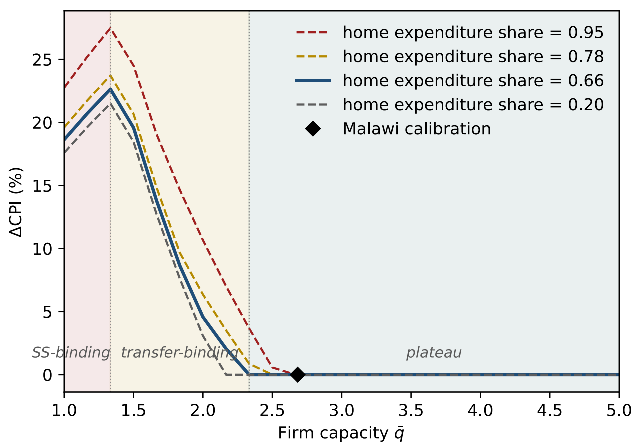

(c) Income–Consumption Multiplier Gap

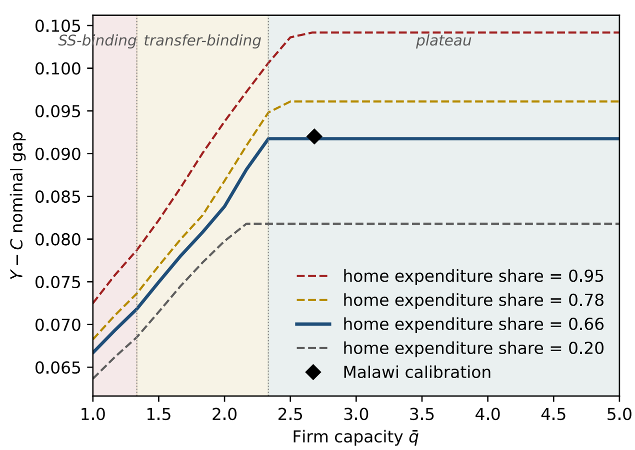

Notes. Each panel plots the cumulative outcome under the experimental transfer (vertical axis) against $\bar{q}$ (horizontal axis) for four values of the home expenditure share (0.20, 0.666, 0.78, 0.95). The Malawi-calibrated home expenditure share is 0.666.

Figure 4: Cost-Effectiveness of (h, v) Allocations.

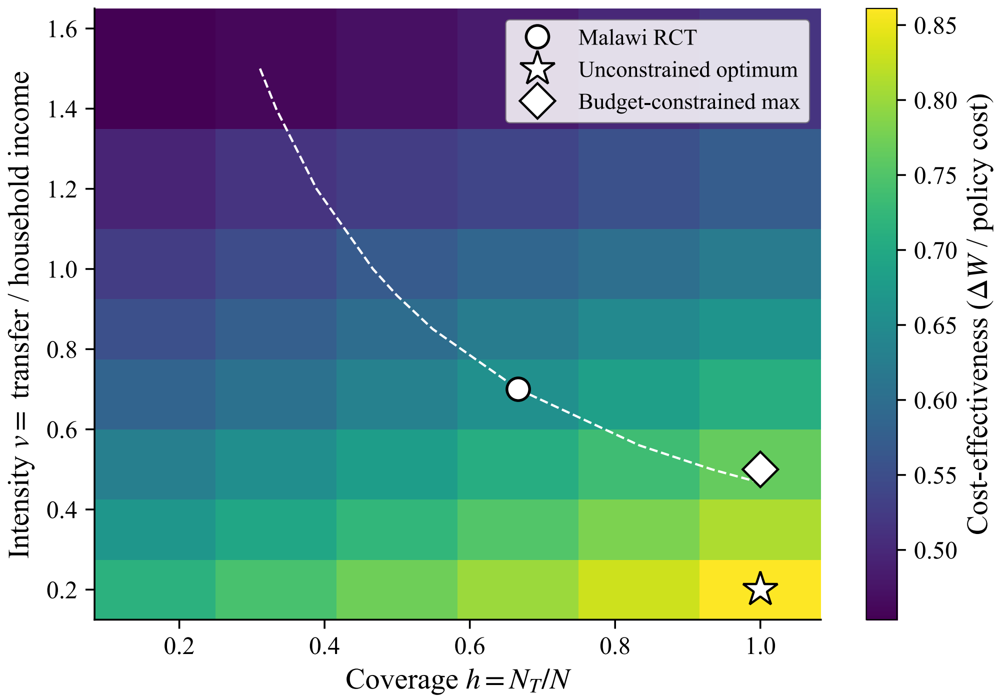

Notes. Malawi high-saturation calibration. Color encodes ΔW/policy cost. The dashed white contour traces iso-budget allocations through the Malawi RCT design at $(h, v) = (2/3, 0.70)$. Moving along the contour from the RCT design toward $(h, v) = (1.0, 0.5)$ raises cost-effectiveness at approximately fixed total cost.

age h, the vertical axis is per-household intensity v, and the color of each cell indicates the welfare gain per unit of policy cost, with brighter (yellow) cells indicating higher cost-effectiveness. The white circle marks the Malawi RCT at $(h, v) = (2/3, 0.70)$; the white diamond marks the budget-constrained welfare maximum at $(h, v) = (1.0, 0.5)$ on the iso-budget contour; the white star marks the unconstrained cost-effectiveness maximum at the bottom-right boundary of the grid. The dashed white contour traces an iso-budget locus, the set of (h, v) allocations whose total transfer cost equals that of the RCT design. Within $\pm 10\%$ of the RCT budget, cost-effectiveness rises as one moves along the contour from the RCT design toward the diamond, which delivers a cumulative welfare gain of 14.33 against the RCT's predicted gain of 11.48.

The 25% higher welfare in the optimal policy reflects two sources. First, additional coverage activates two previously untreated villages whose non-agricultural firms have slack capacity, and each newly treated village adds a nearly independent welfare increment. Second, reducing per-household intensity from 0.70 to 0.50 raises the marginal welfare gain per dollar. Since marginal utility falls as consumption rises, a fixed total dollar amount delivers more aggregate welfare when spread across more recipients than concentrated on fewer. Both effects push in the same direction on the slackness plateau, while neither incurs an inflation penalty. Both the actual RCT design to the full-coverage budget-constrained optimum lies in the zero-CPI region, because per-household intensity (v) in both cases stays below the level at which the transfer saturates slackness (about 0.85 of steady-state income). Therefore, even at full coverage the capacity constraint within any village remains slack.

Above this threshold, the policy ranking reverses. As per-household intensity rises toward 1.5, CPI climbs into single and then double digits and the capacity margin that used be absent reappears, and additional intensity trades welfare for inflation. Full coverage at moderate per-household intensity is therefore the optimal combination of h and v under the Malawi calibration. At the RCT budget, transfers are delivered most efficiently by spreading them across all six villages, and further per-household scaling beyond the saturation threshold trades welfare for CPI inflation.

6 Model Extensions

The baseline calibration of Sections 4 and 5 rests on two modeling choices. Labor reallocates freely across sectors in response to the transfer, and the iceberg trade margin is treated as pure leakage rather than as local non-agricultural income. We relax each in turn here and examine whether the extended model can generate a higher multiplier consistent with the CPI moment. Section 6.1 treats labor-reallocation friction; Section 6.2 treats trade-income pass-through.

6.1 Labor Reallocation Channel

Section 4.3 reported that the calibrated model's free-reallocation GDP multiplier under the Malawi RCT design is 1.83, falling short of Egger et al. (2022)'s point estimate of 2.58. Our model points to one candidate explanation: when households reallocate labor between agriculture and non-agriculture in response to the transfer, an externality among non-agricultural producers depresses the aggregate marginal product of that labor.

The non-agricultural sector in our model produces under a random-shopping technology. A non-agricultural firm's expected revenue rises with the share of time the firm is open, because longer hours raise the probability of being matched with a customer who arrives at a random time. Individuals choose the hours allocated into agricultural and non-agricultural production so that the marginal product of non-agricultural time

Table 4: Model Predictions Under Two Labor-Reallocation Specifications.

Version	λ	$\bar{q}$	GDP mult	ΔCPI (%)	Y−C gap	C mult	Y mult
Free reallocation (main)	1.0	4.82	**1.83**	0.0	0.092	1.60	1.69
Zero reallocation	0.0	3.17	**2.58**	1.85	0.130	2.13	2.26

Notes. Free reallocation ($\lambda = 1$) is the main calibration with full labor mobility between agriculture and non-agriculture. Zero reallocation ($\lambda = 0$) shuts down labor reallocation, holding labor at its steady-state allocation during the transfer. Consumption and income multipliers are nominal; all multipliers are economy-wide.

equals the marginal product of agricultural time. When the transfer hits, demand for non-agricultural goods rises, every operator's expected revenue rises, and operators reallocate time toward non-agricultural production.

Each individual's reallocation, however, lowers the customer-arrival rate per unit of non-agricultural time for every other non-agricultural firm. This externality drives the aggregate non-agricultural marginal product below what each firm perceives privately, so the equilibrium reallocation toward non-agriculture is too high relative to the income-maximizing level. Thus, with labor freely mobile, reallocation lowers total income relative to a counterfactual in which labor is held fixed across sectors. This *competition externality* is one reason behind the gap between the model's free-reallocation multiplier (1.83) and Egger et al. (2022)'s point estimate (2.58).

Table 4 reports two points along this dimension. Under the free-reallocation baseline ($\lambda = 1$, full labor mobility between sectors), the model produces a multiplier of 1.83 with $\bar{q} = 4.82$ identified by the capacity-utilization moment. Re-calibrating the model with $\lambda = 0$ (labor allocation fixed across sectors) and re-identifying $\bar{q}$ against Egger et al. (2022)'s 2.58 yields $\bar{q} = 3.17$; under this re-calibration the model matches the Kenya multiplier exactly. Both values are above the slackness-plateau threshold identified in Section 4.3.

Figure 5 illustrates how the aggregate multiplier depends on the labor-reallocation speed at fixed firm capacity ($\bar{q} = 3.17$). Even a small reallocation friction (about $\lambda = 0.05$) drops the multiplier below 2.4. The income–consumption multiplier gap tracks the same dimension, rising from 0.092 at the free-reallocation end to 0.130 at the fully shut-down end, and provides a second margin on which endline results can shed light.

How big is the reallocation that closes the gap? The reallocation magnitude implied by the free-reallocation baseline ($\lambda = 1$) is quite small: the model predicts roughly a one-hour-per-week shift away from agriculture per treated household. Egger et al. (2022) report

Figure 5: GDP Multiplier as a Function of the Labor-Reallocation Speed λ.

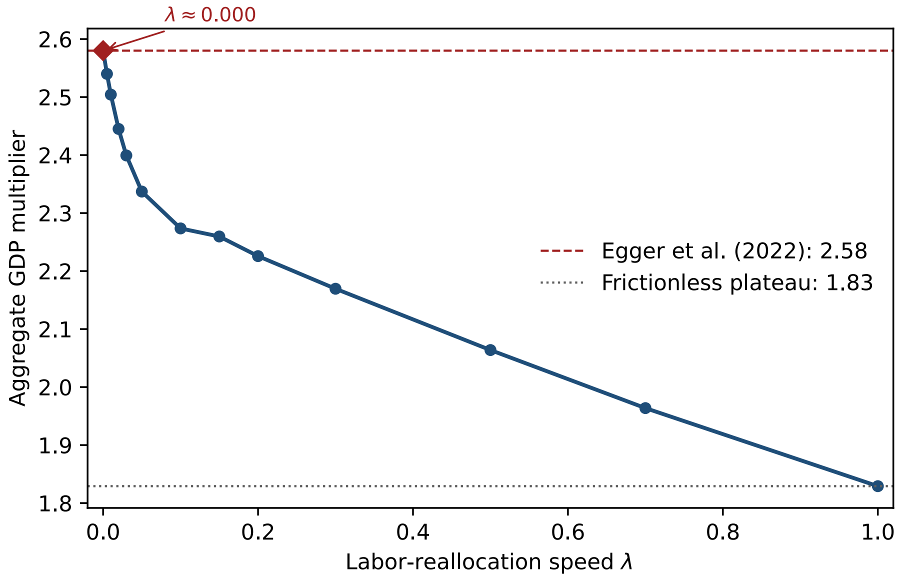

Notes. At fixed $\bar{q} = 3.17$. Horizontal references at Egger et al. (2022)'s point estimate (2.58) and the free-reallocation plateau value (1.83). The crossover near $\lambda = 0$ is steep, so even modest reallocation closes most of the gap.

a small, statistically insignificant point estimate on agricultural labor hours under treatment, with a standard error of 2.3 hours per week. The reallocation channel is consistent with the Kenya RCT, but the experiment lacks the power to identify it. The Malawi endline may be used to test the prediction directly.

If endline GDP, consumption, and income multipliers track 1.83, 1.60, and 1.69, the baseline labor-reallocation specification is consistent with the data. Multipliers that overshoot toward 2.58 with low CPI would prefer a slower reallocation. The reallocation magnitude itself, agricultural-hours change in the treated villages relative to controls, is a direct test of the underlying mechanism.

Figure 6 repeats the cost-effectiveness exercise of Section 5 under zero labor reallocation, with $\bar{q} = 3.17$ recalibrated to match Egger et al. (2022)'s 2.58. The welfare-maximizing allocation remains at $(h, v) = (1, 0.5)$, full coverage at intermediate per-household intensity. The shape of the cost-effectiveness surface is preserved regardless of the introduction of the labor reallocation channel. The coverage-vs-intensity asymmetry that drives Section 5's finding is therefore a feature of the slackness plateau rather than of the reallocation specification.

Figure 6: Cost-Effectiveness of (h, v) Allocations Under Zero Labor Reallocation.

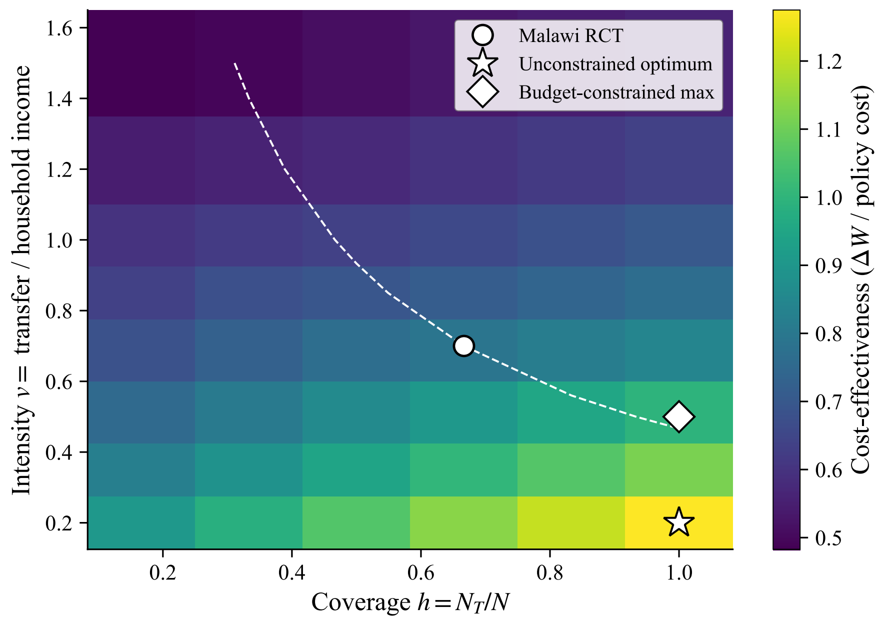

Notes. $\lambda = 0$, with $\bar{q} = 3.17$ recalibrated to match Egger et al. (2022)'s GDP multiplier of 2.58. The welfare-maximizing allocation remains at $(h, v) = (1, 0.5)$, identical to the baseline result of Section 5.

6.2 Trade-Income Pass-Through

The baseline model's iceberg trade cost τ implies that shipping one unit of a non-agricultural good across villages requires more than one unit to be produced, with the extra share melting in transit. The baseline calibration treats the melted share as pure leakage. A natural alternative is to interpret the iceberg margin as transport, distribution, and storage activity that takes place in the exporting village and pays local non-agricultural workers as a byproduct. The trade-income share ϕ governs the fraction of the iceberg margin captured as local non-agricultural income. Pure leakage ($\phi = 0$) is the baseline; full pass-through to the exporting village ($\phi = 1$) is the upper bound.

Table 5 reports the model's predictions under pure leakage and full pass-through, with firm capacity $\bar{q}$ recalibrated at each value to satisfy the capacity-utilization moment. Under full pass-through, the GDP multiplier rises to 2.58 at firm capacity 3.05, closing the entire gap to Egger et al. (2022)'s point estimate. The income–consumption multiplier gap rises to 0.141, and the consumption multiplier rises to 1.893. The mechanism is straightforward. Under full pass-through, demand from treated villages that leaks to neighboring villages now generates non-agricultural income there, which raises aggregate non-agricultural production and hence the multiplier.

Table 5: Sensitivity of Model Predictions to the Trade-Income Share ϕ.

Version	ϕ	$\bar{q}$	GDP mult	ΔCPI (%)	Y$-$C gap	C mult
Baseline ($\phi = 0$)	0.0	4.82	1.83	0.00	0.092	1.60
$\phi = 1$	1.0	3.05	2.58	12.62	0.141	1.89

Notes. ϕ is the share of the iceberg trade margin that returns to the exporting village as local non-agricultural income; $\phi = 0$ is the baseline (iceberg costs are pure leakage); $\phi = 1$ is full pass-through. Each row recalibrates $\bar{q}$ independently to satisfy the capacity-utilization moment. Consumption and income multipliers are nominal and economy-wide.

Full pass-through, however, produces substantial price effects in the treated villages. The same exercise that delivers the 2.58 multiplier produces a 12.62% rise in the treated-village CPI, against Egger et al. (2022)'s observation of zero inflation. The pass-through that closes the multiplier gap also routes a much larger share of demand into binding firm capacity, raising prices. If the Malawi experiment also finds muted price effects, the trade-income channel at this magnitude would therefore be inconsistent with the experimental outcomes, hence should not be viable mechanism behind the experiment.

The exercise narrows the menu of structural channels that can close the multiplier gap without violating the CPI observation. Labor reallocation closes the gap with approximately zero CPI change at $\lambda = 0$ and $\bar{q} = 3.17$. Trade-income closes the gap with 12.62% CPI change at $\phi = 1$ and $\bar{q} = 3.05$. The CPI moment therefore discriminates between the two channels at their gap-closing values, and points toward labor reallocation. A more disaggregated trade-income calibration anchored in the Malawi enterprise census can sharpen the bound; the upper-bound $\phi = 1$ case is sufficient to rule out the channel at this stage.

7 Conclusion

Cash transfers in low-income economies operate through two distinct channels, slackness in non-agricultural production and openness to inter-village trade. We develop a general-equilibrium model that incorporates both channels and calibrate it to baseline data from the Malawi cash-transfer RCT, with one moment from a comparable Kenyan experiment of Egger et al. (2022). The calibrated economy lies on the slackness plateau, where the price response of the slackness channel is fully muted. Within this regime, the welfare-maximizing transfer design spreads transfers across more villages rather than concentrating them at higher per-household amounts. Coverage and intensity are not substitutes on the plateau; each additional treated village adds a nearly independent welfare increment, while pushing per-household intensity beyond the plateau threshold delivers diminishing welfare gains as local prices begin to rise.

The gap between the model's predicted multiplier and the larger one reported by Egger et al. (2022) for a comparable Kenyan transfer experiment admits a candidate structural interpretation in terms of the speed of labor reallocation across sectors. The model points to a competition externality in non-agricultural production under which transfer-induced demand pulls labor toward non-agricultural work at a rate that depresses the aggregate marginal product of that labor. The implied reallocation magnitude is small enough to fall within the standard error band of the Kenyan estimates and is therefore not identifiable on those data. The Malawi endline may be used to measure cross-sector labor reallocation directly and to test this prediction.

The model commits to specific, falsifiable transition-path predictions for the Malawi endline. The income–consumption multiplier gap, read alongside the GDP multiplier and CPI, discriminates among mechanisms; a joint departure of these measures from the calibrated baseline points to which structural force is mis-specified. Two natural extensions of the framework are dynamic transfer designs that exploit the time profile of the multiplier and geographic generalizations that loosen the home-expenditure-share assumption in settings where the openness channel is more likely to bind.

References

Akcigit, Ufuk, Harun Alp, and Michael Peters, "Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries," *American Economic Review*, 2021, *111* (1), 231–275.

Banerjee, Abhijit, Michael Faye, Alan Krueger, Paul Niehaus, and Tavneet Suri, "Universal Basic Income: Short-term Results from a Long-Term Experiment in Kenya," 2023. Unpublished Manuscript, Massachusetts Institute of Technology.

Buera, Francisco J, Joseph P Kaboski, and Robert M Townsend, "From Micro to Macro Development," *Journal of Economic Literature*, 2023, *61* (2), 471–503.

Chen, Chaoran, Diego Restuccia, and Raül Santaeulàlia-Llopis, "Land Misallocation and Productivity," *American Economic Journal: Macroeconomics*, 2023, *15* (2), 441–465.

Cunha, Jesse M, Giacomo De Giorgi, and Seema Jayachandran, "The Price Effects of Cash Versus In-Kind Transfers," *The Review of Economic Studies*, 2019, *86* (1), 240–281.

Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker, "General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya," *Econometrica*, 2022, *90* (6), 2603–2643.

Farhi, Emmanuel and Iván Werning, "Fiscal multipliers: Liquidity traps and currency unions," in "Handbook of macroeconomics," Vol. 2, Elsevier, 2016, pp. 2417–2492.

Fujimoto, Junichi, David Lagakos, and Mitchell VanVuren, "Aggregate and Distributional Effects of 'Free' Secondary Schooling in the Developing World," 2023. Unpublished Manuscript, Boston University.

Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jón Steinsson, "The Slope of the Phillips Curve: Evidence from US States," *The Quarterly Journal of Economics*, 2022, *137* (3), 1299–1344.

Lagakos, David, Ahmed Mushfiq Mobarak, and Michael E Waugh, "The Welfare Effects of Encouraging Rural–Urban Migration," *Econometrica*, 2023, *91* (3), 803–837.

Lewis, W Arthur, "Economic Development with Unlimited Supplies of Labour," *The Manchester School*, 1954, *22* (2), 139–191.

Mendes, Arthur, Wataru Miyamoto, Thuy Lan Nguyen, Steven Pennings, and Leo Feler, "The Macroeconomic Effects of Cash Transfers: Evidence from Brazil," 2023. Federal Reserve Bank of San Fransisco Working Paper 2024-02.

Nakamura, Emi and Jón Steinsson, "Fiscal Stimulus in a Monetary Union: Evidence from US Regions," *American Economic Review*, 2014, *104* (3), 753–792.

Ramey, Valerie A and Sarah Zubairy, "Government Spending Multipliers in Good Times and in Bad: Evidence from US historical Data," *Journal of Political Economy*, 2018, *126* (2), 850–901.

Todd, Petra E and Kenneth I Wolpin, "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American Economic Review*, 2006, *96* (5), 1384–1417.

Walker, Michael W, Nachiket Shah, Edward Miguel, Dennis Egger, Felix Samy Soliman, and Tilman Graff, "Slack and Economic Development," Technical Report, National Bureau of Economic Research 2024.