

Discussion Paper Series – CRC TR 224

Discussion Paper No. 764  
Project A 03

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

Minki Kim<sup>1</sup>  
Mitchell VanVuren<sup>2</sup>

June 2026

<sup>1</sup>University of Mannheim, Email: [minki.kim@uni-mannheim.de](mailto:minki.kim@uni-mannheim.de)

<sup>2</sup>Vanderbilt University, Email: [mitchell.vanvuren@vanderbilt.edu](mailto:mitchell.vanvuren@vanderbilt.edu)

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)  
through CRC TR 224 is gratefully acknowledged.

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

Minki Kim<sup>†</sup>

University of Mannheim

Mitchell VanVuren<sup>‡</sup>

Vanderbilt University

June 2026

## 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 estimate the model on baseline data from an ongoing large-scale cash-transfer experiment in Malawi. The estimated economy lies on a slackness plateau where firms have substantial idle capacity, muting the price response and generating a local GDP multiplier of between 1.1 and 1.5. Within this regime, the welfare-maximizing transfer design spreads transfers across more villages rather than concentrating them at higher per-household amounts. Our approach illustrates the value of pairing a structural model with an ongoing field experiment.

---

\*For helpful comments on this project we thank Dennis Egger (our discussant), Joe Kaboski, and participants of the Carnegie-Rochester-NYU Public Policy Conference on “Macroeconomic Consequences of Micro-Development Policies.” 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.” Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project A03) is gratefully acknowledged. All potential errors are our own.

<sup>†</sup>Email: [minki.kim@uni-mannheim.de](mailto:minki.kim@uni-mannheim.de)

<sup>‡</sup>Email: [mitchell.vanvuren@vanderbilt.edu](mailto: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 productive inputs (Lewis, 1954); labor markets in developing countries may feature perpetually dormant productive capacity due to market failure or other frictions (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 fiscal and monetary 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 in the developing world. 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 and estimate a general equilibrium macroeconomic model that incorporates both slackness and openness in order to quantify and compare the two channels. The model features a flexible local 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. Importantly, the model introduces slackness through a random shopping microfoundation, where firms produce only when they have customers but cannot predict precise arrival times and thus may spend some of their labor time idle and waiting for customers to arrive, which allows for the possibility that productive capacity is underutilized even in steady-state.<sup>1</sup> Finally, there is inter-village trade so that competition *across* villages (i.e. between those receiving and not receiving the transfer) can act to suppress local price increases in response to the transfer, effectively allowing increased local demand to be absorbed via external supply.

Although slackness and openness imply different policy prescriptions, both likely operate to some degree. The key question is their relative importance. The model provides the insight that comparing the sizes of the local consumption and income multipliers in response to the transfer (that is, the multiple of the total transfer amount by which local consumption and incomes increase) is particularly informative in distinguishing between the two channels. The intuition is straightforward — with trade, consumption and income multipliers might not be identical as the transfer can be spent outside the local village. In a world where slackness dominates, local firms can simply activate the dormant production capacity to cater to higher demand triggered by the transfer, increasing both consumption and incomes. In contrast, when openness dominates, the increased demand is met outside the village which, although it contributes to overall production/GDP, leaves *local* incomes relatively unchanged. Thus, given an experimental structure allowing one to measure consumption and income multipliers directly, the size of the gap between the two points directly towards their relative strengths.

---

<sup>1</sup>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.

Such an experiment is currently being conducted, though not yet finished, by GiveDirectly in Malawi. This timing provides a unique opportunity to bring the model to baseline data, use the model to predict key macro-relevant experimental outcomes (e.g. multipliers) and, eventually, assess the successes and failures of the model (including the link between local multipliers and slackness/openness above) once post-treatment data is available. The experiment is randomized at the village level and provides transfers of approximately \$550 per adult (70–100% of average annual household income at baseline) to all adults in treated villages.

We estimate the model using baseline data from this experiment, corresponding to the Chiradzulu region of Malawi. These data are extensive and contain direct measurement of many moments that are highly relevant to our channels of slackness and openness, including measures of capacity utilization and the shares of consumption and intermediate goods that are sourced outside of the local village. As a point of comparison (and to perform some basic model validation), we additionally estimate the model to match the context of a similar experiment in Kenya (Egger et al., 2022).

The estimated model reveals two key insights. The first is that the Malawian economy is both highly slack (due to low capacity utilization) and highly open (due to a large share of intermediate goods being sourced from outside). It rests deep in what we refer to as the “slackness plateau”, where there is enough residual slackness relative to the size of the transfer that productive capacity is *never* exhausted (and additional slackness has no impact, hence the term “plateau”). However, despite being both slack and open, the economy exhibits surprisingly low multipliers: those for overall (i.e. inclusive of spillovers to non-treated villages) GDP, local consumption, and local income are all about 1.1. The reason for this is that the impact of openness changes once sufficiently deep into the plateau; when local production capacity can absorb the entire increase in demand, multipliers depend only on the marginal propensity to spend locally for classical Keynesian reasons, and this is *decreasing* in openness.

The second insight, also related to the surprisingly low multiplier, is that the model predicts a substantial and rapid reallocation of labor towards local non-agriculture in response to the transfer. This is strange because, as noted above, the high level of slackness means that the additional demand from the transfer could be accommodated *without any additional labor input*. Despite this, households shift towards non-agriculture in hopes of capturing a larger portion of the (now larger) customer base. While this works individually, it does not work *in aggregate* as the total size of the customer pool is fixed according to demand. Households generate something akin to a competitive externality on

each other, and this shift in labor results in no additional non-agricultural production in aggregate but reduces agriculture output/income (due to the loss of labor) resulting in substantially smaller multipliers.

This shift occurs and unwinds extremely quickly (non-ag hours increase by 13 percent in the first month after the transfer, and the impact essentially vanishes by month twelve) which may not be realistic. We consider an extended version of the model with a reallocation friction that mechanically reduces the speed at which households can shift labor between sectors.<sup>2</sup> At maximal frictions where reallocation is impossible, the GDP multiplier increases to 1.5, indicating that the impact of this channel is quantitatively substantial. As a point of comparison, the model estimated to match Kenya generates multipliers between 1.6 (no frictions) and 2.5 (maximal frictions). While both are contained in the 95 percent confidence interval, the point estimate from [Egger et al. \(2022\)](#) for Kenya of 2.6 suggests that the frictional case is closer to reality.

Turning to optimality, we use the estimated model to determine how transfers should be optimally structured. In particular, we consider a situation where the total budget for transfers is the same as in the GiveDirectly experiment, but the policy designer can choose between the scope (portion of villages chosen to receive transfers) and size of the transfer subject to this constraint. We consider two notions of optimality. The first is the classic consumption-equivalent welfare criterion while the second is total output (or, equivalently, the size of the GDP multiplier). Surprisingly, these two criteria give opposite conclusions. The GDP-maximizing transfer is large and concentrated in a single village, as concentrated transfers help mitigate the (negative) impact of reallocation described above by minimizing the number of households who are incentivized to shift sectors. The effect is quantitatively small, however, and the welfare-maximizing transfer is as diffuse as possible in order to leverage curvature of the utility function.

Overall, the estimated model predicts multipliers ranging from 1.1 to 1.5, with our preferred estimate closer to 1.5. Malawi appears to be both highly slack and highly open, but behaves more like a slack economy. There is not much gap between the consumption and income multipliers, a key model indicator of slackness versus openness, and optimal policy simulations suggest that welfare is maximized by making sure transfers are sufficiently diffuse.

---

<sup>2</sup>We do not take a stand on whether this friction is technological (i.e. an already planted crop cannot simply be abandoned) or informational (i.e. households do not know how much demand will increase in response to the transfer), though both are likely present to some extent.

**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 a 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.

$$C_{N,v} = \left( \int C_{N,v}(i)^{\frac{\sigma_f-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}} \quad (2)$$

Second, village varieties are aggregated with a constant elasticity of substitution  $\sigma_f$ .

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

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

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

**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 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 ( $a \geq 0$ ) which 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  $T$ .

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

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  $\theta$  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, conse-

quently, severely underpredicts multipliers.

**Non-Agricultural Production:** The novelty of the model lies in the production of the non-agricultural good and draws on Walker et al. (2024), who emphasize the possibility that production is “slack” in developing countries. We operationalize this concept 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.<sup>3</sup> 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  $\omega$  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

---

<sup>3</sup>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.

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.

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

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.<sup>4</sup> 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 found open by a fraction  $l_N(i)$  of buyers.

Upon arrival, a buyer allocates spending across available varieties according to their CES demand described in (2) and subject to the constraint that they can only purchase from firms that are open in the moment that they arrive. 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. Other than whether they are open or closed, firms are symmetric from the perspective of the buyer so that only the total measure of open firms matters. As a result, each buyer 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

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

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

where  $s_{N,v,w}$  is the share of income spent by households in village  $w$  on non-agricultural

---

<sup>4</sup>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.

goods from village  $v$  so that  $S_{N,v}$  is total spending on  $v$ 's non-agricultural goods.<sup>5</sup>

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

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

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

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

The income maximization problem of the household is then given by

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

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.

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

---

<sup>5</sup>Closed-form expressions for these shares in terms of parameters and prices follow from the household's consumption optimization; see Appendix A.1.

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

$$A\gamma l_A^* \gamma^{-1} = (p^* - P_{N,v}^\omega c(z)) \left( p^{*\sigma} \frac{S_{N,v}}{\int l_{N,v}(j) p_{N,v}(j)^{1-\sigma} dj} \right) \quad (14)$$

**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

$$I(y) = (\pi(p_N, q_N) + Al_A^\gamma) y \quad (15)$$

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 Appendix A.2, 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 incomes align with reality.

$$I_v(i) = (\pi(p_N^*, q_N^*) + Al_A^{*\gamma}) y(i) \quad (16)$$

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 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.<sup>6</sup> 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 50 percent of average village income and plot the impact on local non-agriculture prices as well as real GDP (totaling across both the treated and untreated village, i.e. inclusive of spillovers), 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. Non-agricultural prices stay completely flat until the transfer is large enough to overcome underlying slackness, at which point they also begin to appreciate.<sup>7</sup> 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.

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 occurs 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 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

---

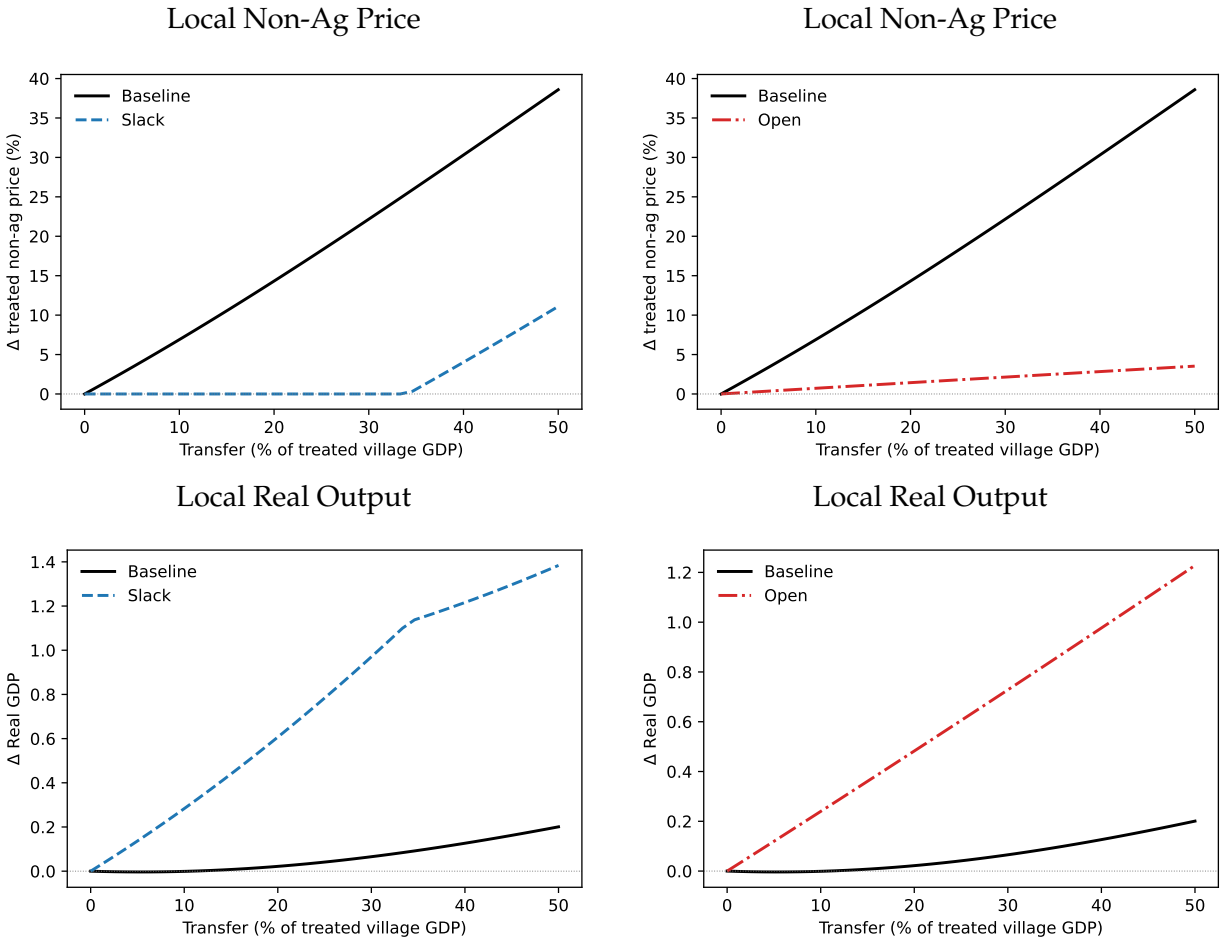
<sup>6</sup>It is important to highlight that these results come from a simplified economy with parameters chosen essentially at random. A full calibrated model is introduced in Section 4.

<sup>7</sup>Recall the agricultural prices are fixed by international markets and thus do not change in any counterfactuals.

Figure 1: Impact of Transfers on Slack vs Open Economies

(a) Slack Economy vs Baseline

(b) Open Economy vs Baseline



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.

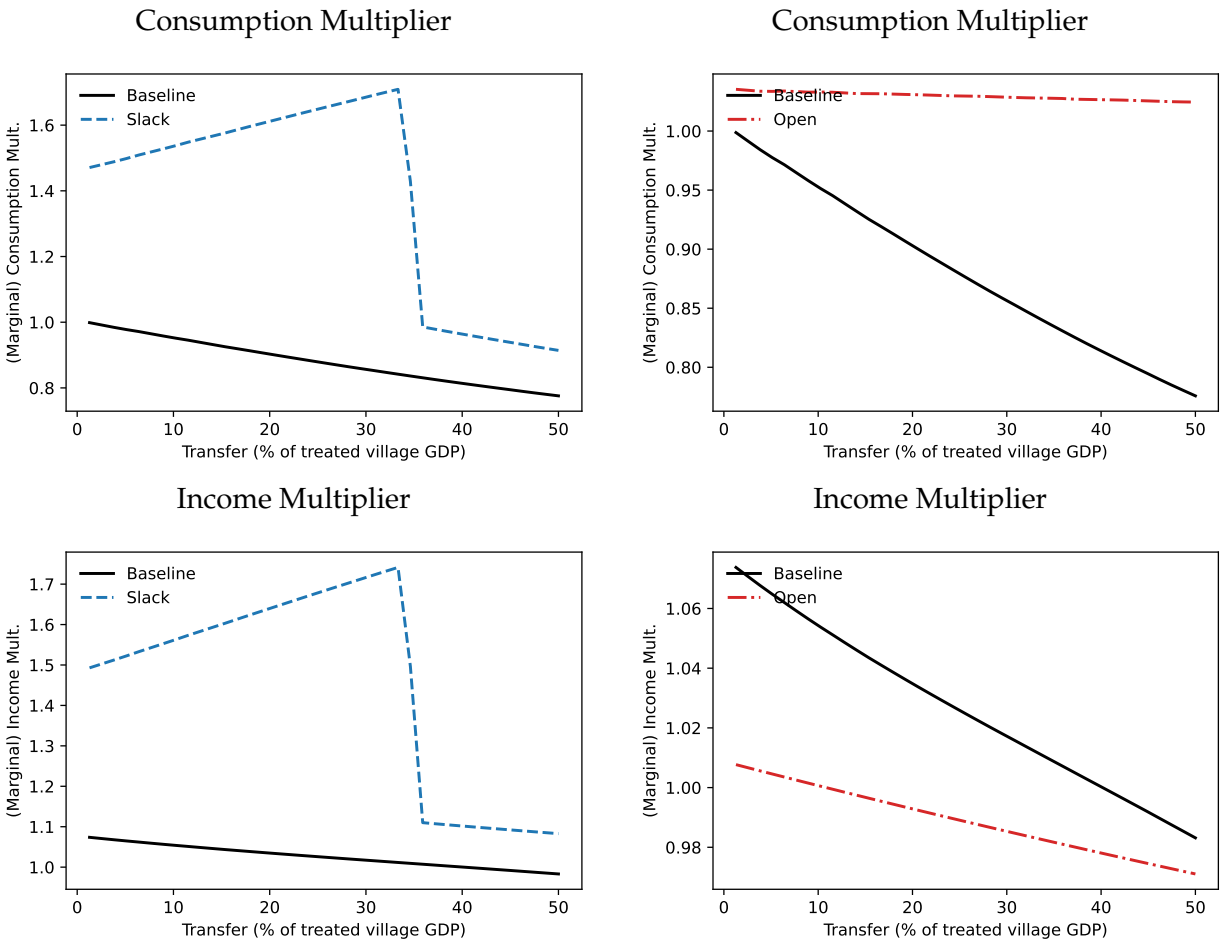
The distinction between the two is seen most clearly by comparing their impacts on local (i.e. exclusive of spillovers) multipliers. We focus in particular on the distinction between consumption multipliers (the multiple by which local consumption increases in response to a transfer) and income multipliers (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

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

(a) Slack Economy vs Baseline

(b) Open Economy vs Baseline



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 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.

## 4 Quantitative Results

Model in hand, our first exercise is to quantify the model using baseline data from the Malawi GiveDirectly cash-transfer experiment. This exercise serves two purposes. First, it allows us to quantitatively examine the mechanisms of slackness and openness in a policy-relevant setting. Second, presenting the model's predictions for various experimental outcomes such as prices, output, and multipliers using baseline data before post-treatment data is available allows us to (eventually) empirically benchmark the model "out-of-sample". Section 4.1 describes the model estimation procedure and targeted moments while Section 4.2 reports the model's baseline predictions for the experimental treatment. Section 4.3 interprets the Malawi calibration in terms of the degree of slackness and openness.

### 4.1 Model Estimation

The parameters to be quantified fall into three categories. The first set is externally calibrated to standard literature values (e.g. discount factors, elasticities, and productivity normalizations). The second set is computed directly from Malawi data (e.g. the agricultural expenditure share, the local input share, and the iceberg trade cost). The third set is jointly estimated via SMM to match five moments — four moments are from the Malawi data while one (the 27-month cumulative MPC out of a cash transfer) is borrowed from

Table 1: Externally Chosen Parameters

Parameter	Description	Value	Source/Target
<i>Normalizations</i>			
$A$	Agricultural TFP	1.0	Normalization
$z$	Non-agricultural TFP	3.92	Normalization
<i>Preferences and technology</i>			
$\beta$	Discount factor	0.95	Standard value
$\rho$	Coefficient of relative risk aversion	1.1	Standard value
$\gamma_A$	Ag returns to scale (DRS)	0.55	Chen et al. (2023)
$\sigma_f$	Across-village CES trade elasticity	4.0	Standard value
$\sigma_{an}$	Ag/non-ag CES elasticity	0.5	Standard value
$y$	Transitory income process	See text	
<i>Data-matched structural parameters/Experimental design</i>			
$\omega$	Intermediate share of local, non-ag	0.54	Enterprise census
$[\tau]$	Iceberg trade cost	3.09	Home share = 0.86
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
$n_{\text{treated}}$	Number of treated villages	4	High-saturation arm

[Egger et al. \(2022\)](#)'s estimates of the Kenya GiveDirectly experiment, since the Malawi experiment is still ongoing at the time of writing.

**Externally Calibrated Parameters:** Table 1 reports the values taken from the literature or computed directly from data. First, TFP in agriculture is normalized to unity and TFP in non-agriculture is normalized to 3.9. This is a strange value, but it has the benefit of ensuring (conditional on other parameters) that the symmetric steady-state equilibrium non-agricultural price is exactly equal to 1, which allows us to accelerate model estimation.

The second block lists preference and technology parameters. We set the discount factor to an annual value of 0.95 while the coefficient of relative risk aversion is set to 1.1. The returns to scale (of labor) in agriculture are taken from [Chen, Restuccia and Santaaulàlia-Llopis \(2023\)](#), who estimate a value of 0.55 for Malawi. The elasticity of substitution between agriculture and non-agriculture is set to 0.5 so that the two are weak complements, consistent with the literature on structural transformation ([Herrendorf, Rogerson and](#)

Valentinyi, 2013; Comin, Lashkari and Mestieri, 2021). The CES demand parameter  $\sigma_f$ , which corresponds to the cross-village trade elasticity, is set to 4.0 — a standard value in the trade literature (Simonovska and Waugh, 2014). Notably, our setting of inter-village trade in a region of Malawi is very different from the international context that is typical for estimates of this parameter; however, lacking any evidence that is directly relevant to our context, we simply use this value.<sup>8</sup>

For the idiosyncratic productivity process, we forgo the typical AR(1) approach and specialize to the binary Markov case. Estimating the model and computing optimal policies involves solving the household’s optimization problem many, many times, and the binary Markov case allows us to write the policy-transition-matrix as a sparse matrix, dramatically accelerating computation and making our more complex quantitative exercise feasible (see Rendahl, 2022, for details). Parameterizing this process, however, turns out to be difficult largely due to the difficulty of precisely measuring household incomes in a rural setting. For conciseness, we constrain this discussion to the appendix, and here simply note that we end up with a symmetric transition matrix with an annual transition probability of 25 percent, normalize the multiplier on income in the high state to unity, and select the multiplier in the low state so that the unconditional variance of log income is 0.5.

**Directly Computed Parameters:** The third block reports parameters computed directly from data in Malawi or corresponding to the experimental design. Of these, the most important are the share of intermediate goods that are sourced locally  $\omega$  and the trade cost matrix  $[\tau]$ . To compute the intermediate goods’ share, we combine data from three different sources. The first is the *household* enterprise survey, which surveys households within villages and reports input usage, divided into intermediate materials/supplies, stock-for-resale, and wages, but not sourcing locations.<sup>9</sup> The second is the *marketplace*

---

<sup>8</sup>We can leverage the fact that the GiveDirectly Malawi data on firms report both sale and source markets to perform a back-of-the-envelope sanity check of this value. We use this data to construct an “extensive margin trade matrix”  $T$  where each entry  $T_{ab}$  is the number of firms who report sourcing a product from market  $a$  and selling it in market  $b$  (done because we lack quantity data). We then estimate a simple gravity equation via PPML:  $T_{ab} = \exp(\gamma_a + \eta_b + \log(d_{ab}))\epsilon_{ab}$  where  $d_{ab}$  is the distance between  $a$  and  $b$  in km. The resulting elasticity is -2.18. Under the assumption that trade costs scale linearly with distance (reasonable in our setting where the mean trade-weighted distance is 7.5km and most transport happens on foot), the implied value for  $\sigma_f$  is  $1 - (-2.18) = 3.18$ . If the elasticity between trade costs and distance is only 0.5 (reflecting, say, the fixed costs of planning), the implied value is  $1 - (-2.18/0.5) = 5.36$ , suggesting our value of 4.0 is reasonable.

<sup>9</sup>From the perspective of the model, which lacks a non-agricultural labor sector, we treat wages paid for labor as an intermediate good and, correspondingly, the household providing the labor as running a non-agricultural firm that produces “labor-services”. This slightly non-standard approach is motivated by the fact that external labor in the sector is extremely rare — wages paid make up less than 1 percent of total input costs — and thus the distinction is quantitatively moot.

enterprise survey, performed at marketplaces (rather than villages). This survey reports the same categories of input usage and also asks about sourcing locations, allowing us to impute the “share spent within the village” for each category in the household survey. Finally, neither survey asks about capital spending which, from the perspective of our model, also serves as an intermediate. Thus we impute both the overall investment shares and “share of investment sourced locally” from [Egger et al. \(2022\)](#), which reports these numbers for Kenya. Aggregating across all categories, we find that 54 percent of intermediate spending is sourced from within the same Traditional Authority as the enterprise, which is our value for  $\omega$ .

Determining the trade costs is more straightforward. For simplicity, we start by assuming that trade costs are symmetric across all foreign villages as well as bilaterally. Villages of course face no trade costs to themselves (i.e.  $\tau_{i,i} = 1$ ), so this essentially leaves us a single parameter to discipline — the iceberg cost of any village buying from any other foreign village  $\tau_{h,f}$ . Combining the symmetry of villages and trade costs with the CES demand structure yields the relationship

$$\text{Share of Spending in Own Market} = \frac{1}{1 + (N - 1)\tau_{h,f}^{1-\sigma_f}}$$

so that conditional on the number of villages  $N$  (discussed below) and the trade elasticity  $\sigma_f$  (discussed above),  $\tau_{h,f}$  exactly determines the share of spending on one’s own market. Using data from the Malawi baseline (and in coordination with [Mendes, Miyamoto, Nguyen and Pennings, 2026](#), who compute an identical number), we compute that 57 percent of all household spending translates to income within a household’s own traditional authority. Given our values for the share of income spent on agriculture (0.28, almost all local), and the share of intermediate goods sourced locally (0.54), matching this value requires that 86 percent of non-agricultural spending occur locally, which corresponds to an iceberg parameter of 3.09.

The remaining parameters all correspond to the particular setting and implementation of the GiveDirectly RCT and are taken from pre-analysis plan documents. The model concept of a village is mapped to a "Traditional Authority" (which is actually a larger designation than a village), as this is the unit of randomization in the trial. Within the trial, there are two arms corresponding to different levels of saturation, with two-thirds of TAs in the high saturation arm receiving treatment while the low saturation arm only treats one-third. We target the high saturation arm and choose  $N = 6$  with 4 TAs receiving

treatment.<sup>10</sup> In treated TAs, 100 percent of households receive an identical transfer equal to 70 percent of household consumption.

**Parameters Estimated via SMM:** Table 2 reports the five SMM-estimated parameters alongside their target moments and the resulting model fit. The steady-state is fairly simple to solve and the mapping between moments and parameters is transparent, so we are able to match all five moments exactly.

The interest rate  $R$  is chosen to match the asset-to-income ratio for Malawi households measured in the baseline survey. This ratio is measured to be about 88 percent, which the model matches with an interest rate just below unity, reflecting a slight negative return on savings. The value of the elasticity of substitution between household-level goods  $\sigma$  is chosen to match a profit-to-revenue ratio of about 31 percent as measured in the enterprise component of the household baseline survey.<sup>11</sup>

The share of households who are hand-to-mouth  $\theta$  is chosen to match the 27 month horizon marginal propensity to consume estimated by Egger et al. (2022) in Kenya. In order to match this, the model needs a HTM share of 88 percent. The relative weight of agricultural consumption in utility,  $\alpha$ , is chosen to match the overall agricultural expenditure share estimated in the baseline household data. We apply a strict definition of agriculture, counting only goods produced/grown by the household itself. For example, maize purchased from a vendor is *not* counted, under the justification that such a transaction actually represents household consumption of a retail service (i.e. the sourcing, transportation, and storage of the maize), while the maize itself is accounted for in the agricultural share of intermediates. Under this definition, we compute that agriculture makes up 28 percent of household consumption.<sup>12</sup>

Finally, the maximum capacity constraint  $\bar{q}$  is chosen to match time-use data from the Malawi enterprise survey. In particular, enumerators were sent to secretly observe businesses and record (in 3 minute intervals), what tasks were being performed. In the base-

---

<sup>10</sup>Under symmetry, the level of  $N$  essentially acts as a normalization. Choosing  $N = 600$  with 400 TAs receiving treatment gives very similar results.

<sup>11</sup>It is worth noting that although the relationship between the markup and  $\sigma$  is determined analytically in the slack regime, the possibility that the capacity constraint binds and leads to a markup that is higher than the usual CES formula necessitates treating  $\sigma$  as an SMM parameter rather than a directly estimated parameter.

<sup>12</sup>Matching a broader definition of agriculture, including purchased foods, requires a substantially more complicated version of the model that draws a distinction between home production and purchased food and accounts for the fact that intermediate usage intensity (i.e. fertilizer) differs substantial between the two. A version of this model (available upon request) generates nearly identical results, motivating our choice to maintain simplicity.

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
$\sigma$	3.25	Profit/revenue ratio	0.308	0.308	HH baseline
$\theta$	0.88	27-month cumulative MPC	0.93	0.93	Egger et al. Table C.I
$\alpha$	0.28	Ag expenditure share	0.276	0.276	HH baseline
$\bar{q}$	8.33	Capacity utilization	0.291	0.291	Enterprise census

line data, 71 percent of time is spent waiting for customers or otherwise idle. Thus we choose  $\bar{q}$  so that average capacity utilization ( $\mathbb{E}(q^*)/\bar{q}$ ) is equal to 29 percent.

**Kenya Estimation:** In order to provide a point of contrast against the model’s predictions for Malawi (as well as some validation), we also estimate a version of the model to match the region of Kenya included in the previous GiveDirectly cash transfer experiment documented in Egger et al. (2022). Given that we lack the raw data from this experiment, we reconstruct the same moments used in the Malawi estimation above as closely as possible using the moments reported in Egger et al. (2022). The details are contained in the appendix. One notable exception is capacity utilization, which was not measured in Kenya. For this, we simply impute the Malawi value.

## 4.2 Running the Experiment in the Model

With the estimated 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 traditional authorities, four treated, full coverage within a treated TA, per-household transfer size 0.70). Thus these headline numbers are essentially the model’s predictions for the experimental outcomes that will eventually be measured and serve as the calibrated model’s pre-registered predictions for the Malawi endline, as they are generated from the model estimated to the baseline data, not fit to the (yet to be observed) endline outcomes.

Table 3 displays the main results, that is, the predicted multipliers and price effects, from this exercise. We note again that there are many different potential multipliers. The first is the aggregate GDP multiplier — that is, the ratio of the increase in total output of all

Table 3: Model Predictions for the Malawi Experiment

Quantity		
	Malawi	Kenya
Aggregate GDP multiplier	1.15	1.60
Consumption mult.	1.08	1.51
Income mult.	1.13	1.58
$\Delta P_{\text{non-ag}}$ (%)	0.0	0.0
$\Delta \text{CPI}_{\text{treated}}$ (%)	0.0	0.0

*Notes.* Cumulative outcomes over the 27-month transition path under the Malawi RCT/estimation (six villages, four treated, within-village coverage of one, per-household transfer size 0.70) and the Kenya RCT/estimation (six villages, four treated, within-village coverage of 0.33, per-household transfer size of 0.75). Consumption and income multipliers are real and compute only within treated villages.

TAs (both treated and untreated) to the total transfers.<sup>13</sup> For Malawi, our baseline model predicts exactly zero inflation and a relatively modest GDP multiplier of 1.15. In essence, there is a sufficiently large combination of slackness and openness to absorb the entirety of the increase in demand with no increases in prices.

The second column displays the same values computed for the version of the model estimated to Kenya. Like Malawi, the transfers generate exactly zero inflation, a result consistent with the experimental estimates. Unlike Malawi, however, Kenya exhibits a higher GDP multiplier of 1.60, which, while within the experimental 95 percent confidence interval, is substantially smaller than the point estimate of 2.58. This raises two questions, which we address in the next two subsections. First, why does the model predict a lower GDP multiplier for Malawi than Kenya? Second, why does the model seem to underpredict the multiplier in Kenya (and what does this imply about the prediction for Malawi)?

### 4.3 The Slackness Plateau

To answer the first question, we need to take a closer look at how slackness and openness interact, and how this interaction applies to the case of Malawi.

We begin by noting that, in terms of slackness, a quantified version of the model can fall into one of three regimes depending on its value of  $\bar{q}$ . The first, which we call the

<sup>13</sup>Notably, this only counts production within the study area (i.e. treated or untreated villages). Hypothetically, if an individual were to receive a transfer and spend it in a town that is not in the study area at all, this would be recorded as consumption but would not show up anywhere in the measured production statistics. Thus “aggregate” here means “at the level of the study” and not “at the level of the country”.

“steady-state binding” regime, is the case where  $\bar{q}$  is low enough that all firms operate at maximum capacity in the pre-transfer steady state (i.e.  $q_{SS}^* = \bar{q}$ ). In this regime, there is no slack in production and any increases in real consumption in response to transfers must come from imports. The second case, which we call the “transfer binding” regime, is the case where  $\bar{q}$  is large enough that there is excess capacity in equilibrium ( $q_{SS}^* < \bar{q}$ ) but small enough that capacity is exhausted for at least one period after the transfer. In this regime, there is some slack that allows for increased production and consumption in response to the transfer but, eventually, this is exhausted and mild price increases are triggered.

The third regime, and the one that turns out to be quantitatively relevant for our estimated model, is what we call the “slackness plateau” (due to the way it looks in the figures below). This is the case where  $\bar{q}$  is large enough that not only is there excess capacity in equilibrium, it is also the case that  $\bar{q}$  never binds in any period after the transfer. In essence, this is the maximally slack case.


Figure 3 demonstrates this by plotting the GDP multiplier, change in CPI, and Income-Consumption Multiplier gap as a function of slackness (governed by  $\bar{q}$  on the horizontal axis) and openness (governed by the home share  $h$  and local intermediate share  $\omega$  and represented by the various lines). The values in our estimated model for Malawi are denoted using a black diamond. The reason for the term “plateau” is clear; at this point (denoted by the green region), further increases in  $\bar{q}$  have exactly zero impact on any baseline or experimental outcomes.<sup>14</sup> The reasoning is clear — this is the point at which there is enough slackness to absorb the entire increase in demand and, as a result, any additional slackness does not matter. Notably, our calibration for Malawi sits deep into this plateau. The estimated value of  $\bar{q}$  is 8.3, whereas the plateau begins at around 3.4. Thus our predictions (at least at current transfer sizes) are very robust to mismeasurement in capacity utilization; the capacity utilization target (that about 70 percent of firm time is spent idle) could be overstated by a factor of two, and our results would not change at all.

More interesting is what the figure reveals about the interaction between slackness and openness. In the steady-state binding region (denoted in red) and most of the transfer-binding region (denoted in yellow), the impact of openness on the multiplier is as described in Section 3. Higher openness increases the multiplier as competition from other villages suppresses price increases. The story changes, however, on the slackness plateau.


<sup>14</sup>To be precise, the baseline level of capacity utilization changes slightly with respect to openness so that each line has a different cutoff separating the regimes. The regions denoted in the figure correspond to the solid blue line that represents the estimated values for Malawi.

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 bias and local intermediate share. The Malawi-calibrated values are (0.57, 0.54).

Here the effect is reversed and the most open economy (dotted gray line) has the lowest multiplier while the most closed economy (dotted red line) has the highest multiplier. This reversal is the result of the fact that, at this point, there is enough slackness that prices do not increase regardless of openness (e.g. subfigure b). In contrast, more openness results in more leakage spending outside the village and thus a lower multiplier. That is, if a larger portion of spending occurs outside the village, only a small portion recirculates as income to other households within the village. In Keynesian terms, under full slackness, the multiplier is determined by the marginal propensity to spend locally which is lower in an open economy.<sup>15</sup>

This fact/reversal is largely responsible for the model's different multiplier prediction between Malawi and Kenya. Recall that the solid blue line represents the parameters estimated for Malawi. At the slackness plateau, the predicted multiplier is 1.1. The values for the home share (95 percent) and the local share of intermediates (91 percent) for the dotted red line were chosen to match the level of openness measured in the alternative model estimated to Kenya (other parameters remain the same as Malawi). For this line, corresponding to the case of Malawi with the level of openness measured in Kenya, the multiplier at the plateau is just north of 1.6. Recall that the model's predicted GDP multiplier for Kenya was 1.6 — from the perspective of the model, this difference in openness *entirely* explains the multiplier difference between the two countries. Both Malawi and Kenya exhibit so much slack that the entirety of the transfer can be absorbed through higher local production, and openness becomes a detriment rather than a boon (at least from a local perspective) as it limits the extent to which this occurs.

#### 4.4 Sectoral Reallocation and Frictions

Our second question was why the model, when estimated to Kenya, seemed to underpredict the multiplier measured in Egger et al. (2022) (1.6 for the model, 2.6 in the data, though with a lower 95 percent bound of 1.34). Understanding this requires looking at sectoral reallocation in the model and, in particular, the shift towards non-agriculture in response to the transfer.

After the transfer occurs, demand for the local good is higher (both absolutely and, due to trade costs, relatively), so households shift their labor away from agriculture and towards the production of non-agriculture. Notably, this shift occurs regardless of the level

---

<sup>15</sup>Technically, spending by one treated village at another treated village would also increase incomes as measured by the experiment, so a more precise accounting would link the multiplier and the marginal propensity to spend *within the study area*.

of slackness in the economy. Though unsurprising at first glance, upon closer inspection this shift is strange. Even under high levels of slackness, where this increase in demand can be accommodated with *no additional labor input*, the shift still occurs. Because output in the non-agricultural sector is fully demand-determined, this increase in labor results in no additional non-agricultural output. The loss of labor in agriculture, however, does reduce output. Thus the net result of this shift is a decline in GDP, which serves to reduce the multiplier.

The effect occurs due to something akin to a competitive externality in the non-agricultural sector. For any individual, increasing their labor input into non-agriculture increases the number of customers that find their enterprise open and, consequently, their sales. However, customers spread their spending across the businesses that they find open and, on the margin, an additional opening results in fewer purchases from each other business. Thus *in aggregate*, additional labor input has essentially no impact on total sales.<sup>16</sup> As noted above, this is intuitive as the case with slackness represents a world where output is entirely demand constrained.

While this reallocation reduces the multiplier, it is not clear that it is a realistic feature of the model. In particular, the model predicts that this reallocation occurs quickly, with non-agricultural hours increasing by 13 percent in the first month after the transfer, and unwinds equally quickly, with hours increased by only 0.2 percent one year after the transfer (and 0.04 percent at the experimental endline). There are many potential frictions preventing such rapid reallocation (one cannot simply abandon an already planted field, information constraints, etc.) and to the extent that reallocation is dampened, multipliers will be higher.


To investigate this effect, we consider an alternative version of the model with a reduced-form reallocation friction.<sup>17</sup> This takes the form of an additional parameter  $\phi$  that represents the portion of the gap between a household's current level of non-agricultural labor input and their desired (optimal) level of non-agricultural labor input that can be closed in a single period. The case of  $\phi = 1$  corresponds to the frictionless model while the case of  $\phi = 0$  corresponds to a model where labor input shares are fixed at their pre-transfer levels and reallocation is shut down altogether.

---

<sup>16</sup>The impact is not precisely zero, as more labor input means that each customer will find more businesses open and CES preferences exhibit love-of-variety, slightly lowering the "effective" price of non-agricultural consumption. However, this effect is small and, under the parameterization that agriculture and non-agriculture are complements, actually magnifies the effect as a lower price of non-agriculture *lowers* the share of income spent on non-ag.

<sup>17</sup>We avoid taking a stance on the nature of the friction, opting instead to simply examine how outcomes change when such a friction is introduced.

Figure 4: Multiplier as a Function of  $\phi$  (Kenya)



*Notes.* This figure plots the long-run cumulative multiplier as a function of the labor adjustment parameter  $\phi$  which governs the speed at which households can shift their labor between agriculture and non-agriculture. The lower line represents the multiplier in the baseline model ( $\phi = 1.0$ ) while the red line represents the point estimate from Egger et al. (2022).

Figure 4 plots the value of the GDP multiplier in the version of the model estimated to Kenya as a function of this friction parameter. The dotted grey line displays the multiplier in the frictionless case, which is about 1.6. The dashed red line corresponds to the point estimate from Egger et al. (2022). From the figure, it is clear that higher frictions in sectoral reallocation (i.e. lower  $\phi$ ) increase the multiplier dramatically. As  $\phi$  approaches zero (the case in which no reallocation occurs) the multiplier increases to 2.5 — essentially in line with the experimental point estimate. Even at  $\phi = 0.5$ , the multiplier is substantially increased to about 2.0.

It is important to note that this result is purely a mechanical outcome of the interactions described above. That is, this increase in the multiplier does *not* occur because frictions somehow amplify the traditional Keynesian channels by inducing individuals to spend more (locally) or because frictions somehow make slackness/openness stronger. Instead, they simply prevent individuals from taking actions in response to the treatment that, in aggregate, lead to lower output and, because the multiplier is experimentally measured by comparing incomes in treated and untreated areas, thus a lower multiplier.

Table 4: Upper- and Lower- Multiplier Bounds (Malawi)

	Lower Bound ( $\phi = 1.0$ )	Upper Bound ( $\phi = 0.0$ )
Predicted Multiplier	1.1	1.5

*Notes.* This table displays the upper and lower bounds of the predicted multiplier in Malawi according to the model with no reallocation frictions ( $\phi = 1.0$ ) and full frictions ( $\phi = 0.0$ ).

**Falsifiability:** In an ideal world, one could directly measure the response of agriculture and non-agriculture hours to the cash transfer as part of the experiment and use this to confirm the model’s prediction that labor shifts away from agriculture as well as measure the magnitude of the shift over time to discipline  $\phi$ . However, in practice, this turns out to be difficult.

First, the Kenyan experiment only measured labor hours at the 27-month endline. Even in the frictionless case, the model predicts that the transfer increases non-ag hours by only 0.4 percent (roughly one minute per week) at this time horizon which would be effectively undetectable even with absurd sample sizes. Second, even the immediate (first month) impacts are small enough that hypothetical tests at this time horizon would be underpowered at sample sizes similar to those in the Kenya experiment.<sup>18</sup> The predicted one-month increase of 13 percent translates to 2.25 labor hours per week. By comparison, the experimental estimate of non-agriculture hours carries a standard error of 2.31 (point estimate 0.62) so that, if the model were exactly correct, a standard two-sided test with  $\alpha = 0.05$  to detect the impact of treatment would have power of only 16 percent. In essence, despite having a very large impact on the predicted multiplier, the quantitative magnitude of this reallocation is effectively unfalsifiable, at least at current sample sizes/experimental setups.

**Application to Malawi:** What does this imply about the model’s predictions for the multiplier in Malawi? Table 4 displays the multiplier value in the case of no reallocation frictions ( $\phi = 1.0$ ) and full frictions ( $\phi = 0.0$ ) which serve as the lower and upper bounds of the multiplier respectively. This represents the range of potential values predicted by the model, without needing to take a quantitative stance on frictions. The range is much smaller than the Kenyan case, covering only 1.1 to 1.5. Though it is noisy, the point estimate of the multiplier from Kenya is clearly more consistent with the highly frictional case. Lacking a better way to pick a particular point out of the plausible range for Malawi,

<sup>18</sup>Egger et al. (2022) report surveying 7,848 households, though the appendix tables do not precisely indicate the sample size used to estimate the treatment effects on labor hours.

we generalize this observation to the Malawian case as well, and 1.5 serves as our pre-registered prediction for the multiplier that will be observed in the experiment.

## 5 Optimal Transfer Design

Given our estimated model for Malawi, and the implied levels of slackness and openness, 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?<sup>19</sup>

In Section 3, we argued that slackness and openness carried different implications for the optimal structure of transfers. Under slackness, the danger was that large transfers to a single village could overwhelm local production capacity, leading to inflation and lower multipliers and, consequently, the optimal transfer structure was small transfers to many villages. In contrast, under openness, the risk was that making transfers across too many villages could endanger the channel of inter-village competition, and optimal transfers were concentrated in the most connected villages.

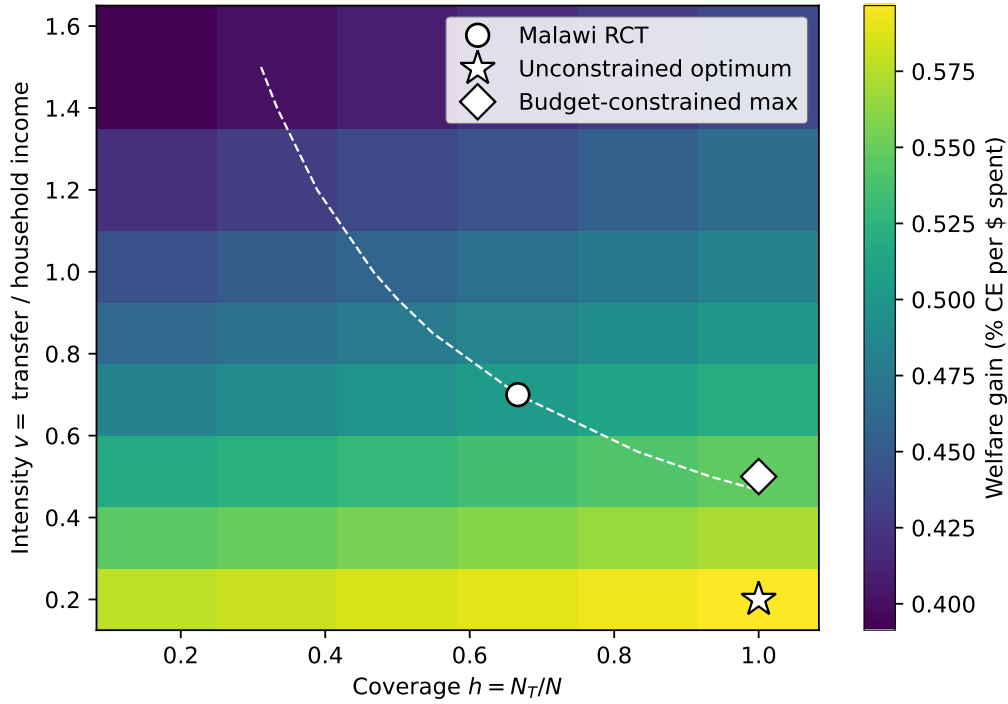
Our estimated model implies that Malawi is both highly slack and relatively open. Thus it is not clear which channel will dominate. To investigate this quantitatively, we set the policy design space as a two-dimensional grid over village coverage  $h$  and per-household transfer intensity  $v$ . The 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. As a benchmark, the Malawi RCT's actual allocation is  $(h, v) = (2/3, 0.70)$ .

We consider optimality with respect to two different outcomes. The first is the natural consumption-equivalent utilitarian welfare criterion. The second is the average income multiplier across the entire study area (i.e. including both treatment and control sites). While the consumption-equivalent approach has the advantage of speaking directly to utility/welfare, curvature in the utility function does bias it towards a “wide” approach to transfers, even if these do not induce the most production. On the other hand, someone

---

<sup>19</sup>Note that the question of optimal *transfers* is different from the question of optimal *policy*. Indeed, the presence of slackness in steady-state and the accompanying competitive externality imply that the presence of a social planner would improve outcomes relative to competitive equilibrium. Interestingly, a planner's optimal allocation would exhibit exactly zero slackness, and any steady-state slackness represents labor that could be productively allocated to agriculture at zero cost. Thus, at least from the perspective of the simple model, a fully optimal policy must involve taxes on the non-agricultural sector or subsidies to agriculture. In this section, we do not consider anything along these lines and instead limit ourselves to investigate how unconditional cash transfers can best be targeted.

Figure 5: Average CE Welfare of  $(h, v)$  Allocations.




Notes. Malawi high-saturation calibration. Color encodes  $\Delta W/\text{policy cost}$ . The dashed white contour traces iso-budget allocations through the Malawi RCT design at  $(h, v) = (2/3, 0.70)$ .

interested in the ability of cash transfers to generate economic growth and production may be more interested in maximizing the average income multiplier, even if that results in lower curvature-inclusive social welfare.

Figure 5 plots average consumption equivalent welfare per dollar spent across the policy grid while Figure 6 plots the GDP multiplier. The horizontal axis is village coverage  $h$ , the vertical axis is per-household intensity  $v$ , and the color of each cell indicates the welfare gain/multiplier, with brighter (yellow) cells indicating higher effectiveness. The white circle marks the Malawi RCT at  $(h, v) = (2/3, 0.70)$ , and the white dotted line represents an “iso-budget” line of  $(h, v)$  combinations that generate the same total cost as the experiment itself. In each figure, the white diamond marks the budget-constrained optimum according to the desired criterion (i.e. the maximizer of the color axis along the iso-budget line). Finally, the white star represents the unconstrained maximizer (in the space of  $(h, v)$  combinations we check). Figure 6 plots the same for the income multiplier, with brighter cells corresponding to larger multipliers.

The welfare results reveal that the welfare-maximizing (per dollar) policy is one that max-

Figure 6: GDP Multiplier of  $(h, v)$  Allocations.



Notes. Malawi high-saturation calibration. Color encodes the GDP multiplier. The dashed white contour traces iso-budget allocations through the Malawi RCT design at  $(h, v) = (2/3, 0.70)$ .

imizes coverage, rather than intensity. That is, more transfers to more households/villages rather than larger transfers to a narrow set of recipients. The difference is significant: the optimal (budget-constrained) policy achieves an increase of about 0.5 percent of welfare per dollar spent, compared to about 0.4 for the least-optimal policy. So welfare is about 25 percent higher under the optimal policy. It is tempting to interpret this result as indicating that slackness dominates openness in terms of optimal policy; recall that under slackness the scale-up risk was that large transfers to few villages may exhaust residual slackness.

The multiplier results in Figure 6, however, call this into question. The multiplier is actually *increasing* in intensity. This effect is so strong that despite also increasing in coverage, the multiplier-maximizing policy actually concentrates all the transfers in a single village at maximum intensity. This effect is not strong enough to overcome the curvature of the utility function, so a utilitarian planner would still disperse the budget as widely as possible, but it does reveal that the outcomes in Figure 5 are purely the result of utility curvature.

Why does the multiplier rise with transfer intensity? Recall from Section 4.4 that the trans-

fer reallocates labor from agriculture toward non-agriculture, and that it is this reallocation that depresses the multiplier. However, the reallocation does not grow one-for-one with the size of the transfer. Each household shifts labor toward non-agriculture up to the point where the marginal gain equals the marginal cost. The marginal gain is roughly constant in the hours shifted, since non-agricultural revenue is linear in the time a firm is open (equation 8). The marginal cost, however, rises steeply because agriculture exhibits diminishing returns (equation 10). Each hour withdrawn raises the marginal product of the labor left in agriculture, and the forgone agricultural output rises convexly. A larger transfer raises non-agricultural demand, and hence the return to reallocating, but the convex cost of withdrawing from agriculture increasingly restrains it; the equilibrium reallocation therefore grows less than proportionally with the transfer. Because the multiplier is output per dollar transferred, a reallocation that grows only sub-linearly imposes a smaller drag per dollar as the transfer grows, leading to a higher multiplier.<sup>20</sup> The effect is bounded: once the transfer is large enough to exhaust even the substantial slackness we estimate for Malawi, local prices begin to rise and the multiplier declines.

## 6 Model Extensions

In this section, we consider two extensions to the baseline model in order to test the robustness of our conclusions to some of the implicit simplifying assumptions. In subsection 6.1, we consider the fact that trade costs may not be purely iceberg and instead represent labor inputs (i.e. the driver of a transport trailer) that accrue as income, increasing the multiplier. In subsection 6.2, we add capital to the household non-agricultural production function, reflecting the fact that transfers may induce households to make additional productive investments into their enterprises which could boost production and thus the observed multiplier.

### 6.1 Trade Costs as Income

The baseline model's iceberg trade cost  $\tau$  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. In this sense, the melted share is treated 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. Producing these services requires local labor,

---

<sup>20</sup>Equivalently, the mechanism implies that with reallocation shut down ( $\phi = 0$ ) the multiplier is invariant to the size of the transfer on the plateau, since it is precisely the reallocation that introduces the dependence.

Table 5: Sensitivity of Model Predictions to the Trade-Income Share  $\gamma$ .

Version	$\bar{q}$	GDP mult	$\Delta\text{CPI}$ (%)	Y–C gap	C mult
Baseline ( $\gamma = 0$ )	8.3	1.1	0.0	0.05	1.1
$\gamma = 1$	9.0	1.5	0.0	0.07	1.4
$\gamma = 1, \phi = 0$	9.0	2.1	0.0	0.10	2.0

*Notes.*  $\gamma$  is the share of the iceberg trade margin that returns to the exporting village as local non-agricultural income;  $\gamma = 0$  is the baseline (iceberg costs are pure leakage);  $\gamma = 1$  represents the case where all trade costs accrue as income. The final row additionally shuts down sectoral reallocation ( $\phi = 0$ ). In the cases with  $\gamma = 1.0$ , we recalibrate  $\bar{q}$  to satisfy the capacity-utilization moment.

and thus under this interpretation, at least a portion of the trade costs should represent additional non-agricultural demand/income and contribute to the multiplier.

To incorporate this in the model, we introduce a new parameter, the trade-income share  $\gamma$ , that governs the fraction of the iceberg margin captured as local non-agricultural income. In order to ship one unit of a non-agriculture good from village  $h$  to village  $v$ , households in  $h$  must produce  $\tau_{h,v}$  units. In the baseline model, the  $\tau_{h,v} - 1$  units that are produced but not consumed simply disappear. In this extension, however, fraction  $\gamma$  (so  $\gamma(\tau_{h,v} - 1)$  units in total) accrue to the exporting village as additional non-ag income. Thus the pure leakage case ( $\gamma = 0$ ) corresponds to the baseline while full capture by the exporting village ( $\gamma = 1$ ) is the upper bound.

Table 5 reports the model’s predictions under pure leakage and full capture. Notably, varying  $\gamma$  does change the amount of non-agriculture production in the pre-transfer steady-state (it is slightly higher due to higher incomes), so we recalibrate  $\bar{q}$  at each value to satisfy the capacity-utilization moment.

The first row repeats the baseline analysis with a GDP multiplier of 1.1. The second row reports the results under full capture. Here, the GDP multiplier rises to 1.5 with the consumption multiplier increasing similarly (the Y-C gap also rises slightly). The mechanism is straightforward. Increasing  $\gamma$  means that a large portion of trade costs recirculate via Keynesian dynamics, rather than simply vanishing from the economy, which increases multipliers for well-understood and classic reasons. More interesting is the fact that the change in the price level continues to be exactly zero. Our estimate for Malawi is far enough into the slackness plateau that even with this higher level of recirculation, the transfer is not enough to overwhelm the underlying slackness.

The final row of the table displays the results when the portion of trade costs captured as income is set to one ( $\gamma = 1$ ) and sectoral reallocation is shut down ( $\phi = 0$ ). Recall that shutting down sectoral reallocation increased the multiplier by 0.4 (from 1.1 to 1.5) and, from the table, we can see that full capture of costs increases it by 0.4 (from 1.1 to 1.5). Making both changes together, however, increases the multiplier by 1.0 to 2.1, more than the sum of the two changes individually. This occurs because of complementarities between the additional income (or, more precisely, a smaller loss of income) from shutting down reallocation and increasing the extent to which trade costs are captured, as the additional income recirculates more effectively.

## 6.2 Enterprise Investment

One notable omission from the model is capital; labor and intermediates (and implicitly land) are the only factors of production. This is motivated largely by the setting as well as existing evidence on cash transfers. First, household non-agricultural production in the experimental setting of Malawi is essentially capital-free. Only 3.3 percent of such enterprises report “Own[ing] any machines, tools, or electronics.” Second, the most comparable experimental evaluation of transfers (Egger et al., 2022) in Kenya finds no impact of transfers on investment or capital (with, possibly, the exception of inventories which while technically investment map better to our model concept of intermediates). Still, it is possible that transfers allow/induce enterprises to increase productivity by investing in capital and, if this is the case, this will impact multipliers. In this section, we extend the model to incorporate a simple concept of capital and test the robustness of our conclusions to this.

We introduce a discrete choice for households: whether or not to purchase a machine. Households that own a machine experience higher non-agricultural productivity while households without a machine decide each period whether or not they would like to purchase one. This makes machine ownership as a new discrete state variable of the model. Consistent with the experimental context, machines are sourced from outside the study area and at a fixed price. Machines decay with some probability at the end of every period, at which point a machine-owning household becomes a machine-free household.

This extension introduces some new parameters, which we do our best to discipline. Notably, the discrete concept of machine ownership maps neatly into the survey question referenced above about whether an enterprise owns any machines, tools, or electronics, all of which we consider a “machine”. Unfortunately, the data here are quite noisy (and cross-sectional, rather than experimental) and, if anything, indicate that machine-owning

enterprises are *less* productive and smaller than the machine-free, even after a fair number of controls.

We opt to handle this by exogenously setting the productivity benefit of owning a machine to 50 percent, and then calibrating the cost of the machine to match the observed adoption rate of 3.3 percent. In essence, this is akin to specifying the machine as a lumpy investment with some rate of return and choosing the extent of “lumpiness” exogenously while allowing adoption to discipline the return. Finally, lacking any time-series data in our setting, we set the depreciation rate of the machine to 50 percent annually, reflecting the high levels of depreciation in developing settings (e.g. [Graff 2026](#)).<sup>21</sup> After introducing this extension, we perform an entire model re-estimation, so that this extended model matches the main moments just as well as the baseline model.

The estimated cost of the machine is about 52 percent of annual income, which corresponds to a net rate of return of about 10 percent annually. Our assumption of a large productivity benefit (which necessitates a large cost to match adoption rates) translates to a very lumpy investment. We view this as giving this channel the best possible chance to have a large impact, as a lumpier investment increases the impact of the transfer, which is about 70 percent of household income and thus large enough to finance the machine, by making it more likely for households to be “pushed over the edge”.

In this extended model, the GDP multiplier is higher, though only marginally. It increases from 1.1 in the baseline model to 1.2 in the extended model. The portion of households that own a machine increases in response to the transfer and peaks at 4.4 percent, compared to 3.3 percent baseline. The reason for these muted effects is two-fold. First, the model is estimated to match an observed MPC of 0.88. Conditional on matching this moment, only a small portion of the transfer is available for investment. Second, the estimated model featured extensive slackness, which means enterprises are not constrained in their ability to produce enough to meet demand but are instead constrained in terms of demand. While investing in a machine mitigates this to some extent by allowing one to lower prices and compete for customers, the core issue facing firms is not one of productive capacity. Overall, we interpret these results as providing strong evidence that, even under relatively favorable assumptions, allowing for investment/capital does not substantially alter our quantitative conclusions.

---

<sup>21</sup>We note here that we have performed this exercise with a wide variety of values of depreciation, and it seems to make no substantial difference.

## 7 Conclusion

Cash transfers in low-income settings raise consumption and incomes by more than their face value while leaving local prices largely undisturbed. Two explanations for this pattern, slackness in local production and openness to trade with other markets, carry sharply different implications for how transfers ought to be designed, yet both are plausible and both likely operate to some degree. We build a general-equilibrium model that nests the two channels and estimate it on baseline data from the ongoing GiveDirectly experiment in Malawi, with a parallel estimation to the Kenyan experiment of [Egger et al. \(2022\)](#) as a point of comparison. Our model predicts that the Malawi economy is both highly slack and highly open, yet it behaves like a slack economy: transfers generate no measurable inflation and an overall multiplier of between 1.1 and 1.5.

An important lesson from the exercise is that slackness and openness are not two interchangeable ways of flattening the local supply curve, but rather interact in a non-trivial way. The calibrated model places the Malawi economy on what we refer to as the “slackness plateau”, where there is so much slackness relative to the size of the transfer that productive capacity is never exhausted. Away from the plateau, openness raises the multiplier because competition from other villages holds down local prices and lets the real transfer go further. But once an economy is on the slackness plateau, prices stop responding regardless of openness and all that openness does is determine how much of each transferred dollar leaks out of the village. In this regime a more open economy has a *lower* local multiplier, for the most classical of Keynesian reasons. This is why Malawi’s predicted multiplier sits below Kenya’s. Both economies exhibit high slackness, but the Malawi experiment areas are much more open than their Kenyan counterparts.

The multiplier is also shaped by how a transfer reallocates labor across sectors. In the model, higher demand for non-agricultural goods makes non-agricultural sector more attractive, so households shift labor out of agriculture, hoping to capture more of the new spending. However, customers spread their spending across the businesses that they find open and, on the margin, an additional opening results in fewer purchases from each other business. The reallocation therefore pulls the multiplier down as the labor input drained from agriculture lower agricultural output. How far the multiplier falls depends entirely on how quickly labor can move between sectors, and shutting down this reallocation entirely raises the multiplier from 1.1 to 1.5.

Lastly, asked how best to spend a fixed transfer budget, the optimal policy depends on what is being maximized. If the objective is welfare, transfers should be spread as widely

and thinly as possible, because the curvature of utility rewards reaching more households. If the objective is the average income multiplier, transfers should instead be concentrated, because a larger per-village transfer dampens the reallocation externality that depresses production.

At the time of writing, the Malawi experiment is still in the field, so every number here is a prediction rather than a fit. The estimated model commits to a multiplier near 1.5 and muted prices, and the endline data will either bear these out or not. We see this as the central payoff of pairing a structural model with an experiment in progress: the model is disciplined out of sample and makes falsifiable predictions for the experiment. Because it commits to a whole set of predictions rather than a single number, the ways in which the endline departs from them (whether across the board or only on particular margins) will help locate where the model most needs sharpening. The same exercise points to what future experiments should measure. For instance, the speed of sectoral reallocation, the mechanism that matters most for the multiplier, plays out at high frequency and early in the transition. This is exactly what experiments of the current size are least able to measure. Future evaluations that follow households at higher frequency, and earlier in the transition, would better placed to detect it.

## 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 Santaeuilàlia-Llopis**, “Land Misallocation and Productivity,” *American Economic Journal: Macroeconomics*, 2023, 15 (2), 441–465.
- Comin, Diego, Danial Lashkari, and Martí Mestieri**, “Structural change with long-run income and price effects,” *Econometrica*, 2021, 89 (1), 311–374.
- 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.
- Graff, Tilman**, “Depreciation and Growth: Evidence from Machine Repair in Uganda,” 2026.
- 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.
- Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi**, “Two perspectives on preferences and structural transformation,” *American Economic Review*, 2013, 103 (7), 2752–2789.

- 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, and Steven Pennings**, “Cash Transfers in Developing Countries: Local Multipliers, Aggregate Effects, and External Adjustment,” Technical Report, CRNYU 2026.
- , –, –, –, –, and **Leo Feler**, “The Macroeconomic Effects of Cash Transfers: Evidence from Brazil,” 2023. Federal Reserve Bank of San Francisco 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.
- Rendahl, Pontus**, “Continuous vs. discrete time: Some computational insights,” *Journal of Economic Dynamics and Control*, 2022, 144, 104522.
- Simonovska, Ina and Michael E Waugh**, “The elasticity of trade: Estimates and evidence,” *Journal of international Economics*, 2014, 92 (1), 34–50.
- 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.

# A Model Appendix

## A.1 Derivation of the Household Demand System

A household resident in village  $w$  has period- $t$  resources  $Ra + I_w + T_w$  from the budget constraint (5) and divides them between savings  $a'$  and total consumption expenditure  $E_w \equiv Ra + I_w + T_w - a'$ . The household then allocates  $E_w$  across the agricultural good and the non-agricultural varieties produced in all villages.

**Within-village allocation.** A buyer arriving in the village- $v$  non-agricultural market at a random time encounters the subset of firms that are open at the moment of arrival. Because each firm  $j$  is open with probability  $l_{N,v}(j)$  independently of the arrival time, the buyer's expenditure share on firm  $i$ , conditional on  $i$  being open, is the standard CES share evaluated against the integral of opening-weighted prices:

$$\frac{p_{N,v}(i)^{1-\sigma}}{\int l_{N,v}(j) p_{N,v}(j)^{1-\sigma} dj} \quad (\text{A1})$$

Multiplying by the open-probability  $l_{N,v}(i)$  gives the unconditional expected expenditure share on firm  $i$ :

$$\omega_{N,v}(i) = \frac{l_{N,v}(i) p_{N,v}(i)^{1-\sigma}}{P_{N,v}^{1-\sigma}}, \quad P_{N,v} \equiv \left( \int l_{N,v}(j) p_{N,v}(j)^{1-\sigma} dj \right)^{\frac{1}{1-\sigma}}. \quad (\text{A2})$$

The price index  $P_{N,v}$  is the cost of one unit of the village- $v$  non-agricultural aggregate. It differs from the standard CES price index by the opening-hours weight  $l_{N,v}(j)$ , which reflects the random-shopping friction. This is the same  $P_{N,v}$  that enters the firm's marginal cost in Section 2 via the Cobb-Douglas intermediate input bundle.

**Cross-village allocation.** Given the village-level price indices and the iceberg cost  $\tau_{w,v}$ , the delivered price of village  $v$ 's non-agricultural aggregate to a buyer in  $w$  is  $\tau_{w,v} P_{N,v}$ . The CES aggregator (3) over villages yields

$$\pi_{N,v}^{(w)} = \left( \frac{\tau_{w,v} P_{N,v}}{P_N^{(w)}} \right)^{1-\sigma_f}, \quad P_N^{(w)} \equiv \left( \sum_{v'=1}^N (\tau_{w,v'} P_{N,v'})^{1-\sigma_f} \right)^{\frac{1}{1-\sigma_f}} \quad (\text{A3})$$

where  $\pi_{N,v}^{(w)}$  is the share of household  $w$ 's non-agricultural expenditure that goes to village  $v$ 's goods, and  $P_N^{(w)}$  is the cross-village non-agricultural price index faced by a buyer in  $w$ .

**Top-level allocation.** The aggregator (4) splits expenditure between agriculture (price normalized to one) and the non-agricultural aggregate (price  $P_N^{(w)}$ ). Cost minimization yields the household's non-agricultural expenditure share and overall consumer price index:

$$\xi_N^{(w)} = (1 - \alpha) \left( \frac{P_N^{(w)}}{P^{(w)}} \right)^{1 - \sigma_a}, \quad P^{(w)} \equiv \left[ \alpha + (1 - \alpha) (P_N^{(w)})^{1 - \sigma_a} \right]^{\frac{1}{1 - \sigma_a}}. \quad (\text{A4})$$

The agricultural expenditure share is  $\xi_A^{(w)} = 1 - \xi_N^{(w)} = \alpha (P^{(w)})^{\sigma_a - 1}$ .

**Income share.** The product  $\xi_N^{(w)} \pi_{N,v}^{(w)}$  is the share of household  $w$ 's consumption expenditure on village  $v$ 's non-agricultural goods. The corresponding income share multiplies it by the consumption-to-income ratio

$$\eta_w \equiv \frac{E_w}{I_w} = \frac{Ra + I_w + T_w - a'}{I_w}, \quad (\text{A5})$$

giving

$$s_{N,v,w} = \eta_w \cdot \xi_N^{(w)} \cdot \pi_{N,v}^{(w)} = \eta_w \cdot (1 - \alpha) \left( \frac{P_N^{(w)}}{P^{(w)}} \right)^{1 - \sigma_a} \left( \frac{\tau_{w,v} P_{N,v}}{P_N^{(w)}} \right)^{1 - \sigma_f}. \quad (\text{A6})$$

For the  $\theta$ -share of hand-to-mouth households,  $a' = 0$  and  $\eta_w = (Ra + I_w + T_w)/I_w$  in closed form. For the remaining  $(1 - \theta)$  optimizing households,  $a'$  is the saving policy from the recursive household problem in Appendix A.2. All other terms in (A6) depend only on parameters and on the equilibrium price-and-hours array  $\{p_{N,v}(j), l_{N,v}(j)\}$ .

**Firm demand curve.** Substituting (A6) into (9) yields  $S_{N,v}$ . Firm  $i$ 's expected revenue is  $\omega_{N,v}(i) S_{N,v}$ , and dividing by  $p_{N,v}(i)$  gives quantity sold:

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

which is equation (8) in Section 2.

## A.2 Recursive Competitive Equilibrium

Each household has permanent type  $(v, z, \chi)$ , where  $v$  indexes the village,  $z \sim G$  is non-agricultural productivity, and  $\chi \in \{0, 1\}$  flags whether the household is hand-to-mouth

(probability  $\theta$ ). The endogenous state is  $(a, y)$  with  $a \geq 0$  and  $y \in \{y_l, 1\}$  following the Markov matrix  $M$ . Let  $\mu_v(a, y; z, \chi)$  be the cross-sectional distribution within village  $v$ .

A recursive competitive equilibrium is a tuple

$$\{V_\chi, g_\chi, p_{N,v}^*(z), q_{N,v}^*(z), l_{N,v}^*(z), l_{A,v}^*(z), P_{N,v}, P_N^{(w)}, P^{(w)}, S_{N,v}, \mu_v\}$$

satisfying the following conditions for every village  $v$  and household type  $(z, \chi)$ .

1. **Household optimization.** For optimizing households ( $\chi = 0$ ) in village  $v$ , the value function and saving rule solve

$$V_0(a, y; v, z) = \max_{a' \geq 0} \frac{(C^*)^{1-\rho}}{1-\rho} + \beta \mathbb{E}[V_0(a', y'; v, z) | y],$$

where  $C^*$  is the CES aggregate (2)–(4) at total consumption expenditure  $E = Ra + I_v(z, y) + T_v - a'$ , with cross-good shares from Appendix A.1. The saving rule  $g_0(a, y; v, z)$  is the maximizer. Hand-to-mouth households ( $\chi = 1$ ) take  $g_1 \equiv 0$ , with  $C^*$  determined by the same CES allocation at  $E = Ra + I_v(z, y) + T_v$ .

2. **Income maximization.** The static policies  $(p_{N,v}^*(z), q_{N,v}^*(z), l_{N,v}^*(z), l_{A,v}^*(z))$  solve the income-maximization problem (6)–(14) under the time constraint (11). Realized income is  $I_v(z, y) = [\pi(p_{N,v}^*, q_{N,v}^*) + A(l_{A,v}^*)^\gamma] y$ .
3. **Aggregate consistency.** The price indices  $P_{N,v}, P_N^{(w)}$ , and  $P^{(w)}$  are defined in (A2)–(A4). Total village- $v$  non-agricultural spending is

$$S_{N,v} = \sum_{w=1}^N \int s_{N,v,w} I_w(z, y) d\mu_w(a, y; z, \chi),$$

with  $s_{N,v,w}$  given by (A6).

4. **Market clearing.** Non-agricultural variety markets clear through the demand curve (8), enforced by the income-consistency condition (16). The asset market clears at the exogenous return  $R$ . Agricultural goods are freely traded internationally at unit price.
5. **Stationarity.** The distribution  $\mu_v$  is invariant under  $g_\chi$  and  $M$ . Transition paths under time-varying transfer profiles  $\{T_v(t)\}$  are sequences of intra-period equilibria with  $\mu_v$  updated each period by one-period-ahead application of  $g_\chi$  and  $M$ .

## B Data Appendix

### B.1 Parameterizing the Income Process

The largest difficulty in parameterizing the income process is the lack of precise, panel measurement of incomes in rural African settings in general and Malawi in particular. We rely on De Magalhães and Santaaulàlia-Llopis (2018) who use World Bank Living Standards and Measurement Survey (LSMS) data from Malawi, Uganda, and Tanzania to estimate various properties of income process, as well as our own computations using rural households from the Malawi LSMS.

Our first approach is to estimate an AR(1) income process directly on all four waves of the Malawi LSMS. We do this at the household level using “income per capita” as our outcome variable to control for household size. We use the Arellano-Bond estimator to control for household fixed effects so that the estimation picks up only the transitory component of income, corresponding to  $y$  in the model. Our estimated persistence parameter is 0.125 and the implied variance of the innovation is 1.326. Notably, there are three years between each wave of the Malawi LSMS, so the implied annual persistence is 0.50 ( $0.125^3$ ) and the variance of an annual innovation could be approximated as  $1.01 \left( \frac{1.326}{1 + 0.125^{2/3} + 0.125^{4/3}} \right)$ .

At first glance, this seems satisfactory. For a symmetric transition matrix, the autocorrelation between outcomes is given by  $2p_{stay} - 1$ , so we could choose a staying probability of 0.75 (transition probability of 0.25) and select the ratio between the high and low state to match an implied unconditional log income variance of 1.347 ( $\frac{1.01}{1 - 0.5^2}$ ). The issue is that this variance is implausibly large. For example, the implied Gini coefficient (under a log-normal income distribution) is given by  $2\Phi\left(\frac{1.347^{0.5}}{\sqrt{2}}\right) - 1 = 0.59$  for *only* the transitory component. The World Bank reports a Gini coefficient of about 0.39 for Malawi, so this is clearly implausibly large.

The straightforward solution to this puzzle is that the transitory variance from the LSMS data is plagued by measurement error, and the value reflects both actual variation in underlying income as well as substantial noise in the estimation process. De Magalhães and Santaaulàlia-Llopis (2018) estimate a similarly large number (1.09) across all three of Malawi, Uganda, and Tanzania, suggesting that this is not a problem unique to Malawi. However, without the ability to detect measurement error directly, we are unable to determine what portion of our value of 1.347 represents true variation.

We opt to use the aggregate Gini coefficient to construct an upper bound on the variance. In particular, the World Bank’s reported Gini value of 0.39 corresponds to a log variance

of 0.5, and we pick the relative incomes of the high and low state to match this value. This is a upper bound in the sense that we are attributing the entirety of income difference within the country to transitory shocks, which is almost certainly not true. De Magalhães and Santaaulàlia-Llopis (2018) do, however, document extremely high levels of income mobility. For example, over 5 percent of households in the top quintile of income fall to the bottom quintile within two years (compared to 0.2 percent in the US), suggesting that this approximation may be more correct than incorrect.