

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

Minki Kim[†]

University of Mannheim

Mitchell VanVuren[‡]

Vanderbilt University

September 2026

Abstract

This paper develops a general equilibrium model to rationalize why large-scale cash transfers in low-income settings generate high multipliers with little movement in prices on average. We provide two mechanisms, slackness in local production and openness to trade with other villages, and show that their interaction determines the multiplier. We estimate the model on baseline data from an ongoing cash-transfer experiment in Malawi. The estimated economy lies on a slackness plateau, where idle capacity absorbs the transfer without price increases while openness lowers the multiplier. Our preferred model specification predicts a study-area income multiplier of around 1.4. 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 part of 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.

[†]Email: minki.kim@uni-mannheim.de

[‡]Email: mitchell.vanvuren@vanderbilt.edu

1 Introduction

Cash transfers have become one of the most widely used social safety net programs in the developing world. Understanding the general equilibrium effects of these programs is especially important in low-income settings, where large-scale transfers impose significant fiscal burdens. A growing body of empirical work has examined this question through large-scale randomized controlled trials. What is surprising is that these studies find little movement in prices on average while finding positive local transfer multiplier effects, with total consumption and income increasing 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 “slackness” in the utilization of productive inputs (Lewis, 1954): labor markets in developing countries may feature perpetually dormant productive capacity in equilibrium because of market failure or other frictions (Walker, Shah, Miguel, Egger, Soliman and Graff, 2024). In this case, transfers can activate the 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 (Nakamura and Steinsson, 2014; Farhi and Werning, 2016). In a village that trades extensively with external markets, transfer-induced demand can be met through increased consumption of goods from outside the village, as external markets provide an additional source of supply while local prices barely move.

In this paper, we develop and estimate a general equilibrium model that incorporates both slackness and openness. The model features a flexible local price level, so that it can accommodate the impact of cash transfers on prices. It also allows the extreme potential outcome in which additional money translates into higher prices without any increase in production. Slackness is incorporated through a random-shopping microfoundation. In the model, firms produce only when they have customers but cannot predict precise arrival times, and thus may spend some of their labor time idle waiting. As a result, productive capacity can be underutilized even in steady state.¹ Openness enters through inter-village trade with iceberg costs. Trade between villages (between those receiving

¹Imagine, 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.

and not receiving the transfer) can suppress local price increases, as external supply absorbs part of the increase in local demand.

The model implies that slackness and openness are not rival explanations, but mechanisms whose interaction determines both price responses and the size of the multiplier. A large multiplier with little price response arises when households spend nearly all of the transfer, productive capacity is sufficiently slack to accommodate the resulting demand, and the economy is closed enough that spending stays local. When slackness is high, prices are flat and the multiplier is governed by the propensity to consume locally. In such a case, greater openness lowers the multiplier by allowing part of each transferred dollar to leak out of the village. When slackness is low, however, openness plays the opposite role: competition from other villages suppresses local price increases and thereby raises the real multiplier.

How does a researcher tell which force is quantitatively more relevant in a given setting? The model shows that the relative sizes of the local consumption and income multipliers (that is, the multiple of the total transfer amount by which local consumption and incomes increase) are particularly informative for distinguishing between the two channels. With trade, consumption and income multipliers might not be identical as the transfer can be spent outside the village. When slackness dominates, local firms can simply activate the dormant productive capacity to meet the increased demand, raising both local consumption and local income. In contrast, when openness dominates, more of the additional demand is met outside the village, which raises consumption without generating a corresponding increase in *local* income. The gap between the two multipliers therefore reveals how much of the transfer-induced demand is met locally rather than leaking out.

We estimate the model using baseline data from an ongoing experiment by GiveDirectly in Malawi. The experiment provides transfers of approximately \$550 per adult (70–100% of average annual household income at baseline) to all eligible adults in treated villages, which are selected through cluster randomization.² These data are extensive and contain direct measures of many moments that are relevant to our channels of slackness and openness, including capacity utilization and the shares of consumption and intermediate goods that are sourced outside the local villages. Moreover, the ongoing nature of the experiment provides a unique opportunity to use the model to predict key macro-relevant experimental outcomes (e.g. multipliers), and later assess the successes and failures of the

²The study area, the Chiradzulu district of Malawi, comprises ten Traditional Authorities (TAs), 113 Group Village Headman (GVH) areas, and 920 villages. Each TA is randomly allocated a level of saturation (the share of its GVH areas that are treated), and within treated GVH areas, two-thirds of villages are treated.

model once post-treatment data are available. As a point of comparison and for a basic validation exercise, we additionally estimate the model to match the multipliers estimated from a similar experiment in Kenya ([Egger et al., 2022](#)).

The estimated model suggests that the Malawian experiment study area is both highly slack, with 71 percent of enterprise time spent idle, and highly open, with 60 percent of intermediate inputs and 30 percent of household spending sourced outside the home area. The estimated economy is 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* fully exhausted. Despite being both slack and open, the economy exhibits a surprisingly modest multiplier. Our preferred model specification predicts an income multiplier of 1.4 inclusive of spillovers to untreated villages, with a consumption multiplier of about the same size and no movement in prices. The reason is as described above: high openness dampens the multipliers when local capacity can actually absorb the entire increase in demand. Because the Malawian economy we consider is highly slack, its high openness decreases the multipliers. The model implies that this is the key mechanism that accounts for the difference between Malawi and Kenya. The same model estimated to the Kenyan experiment of [Egger et al. \(2022\)](#) predicts a multiplier of 2.0, compared to the experimental point estimate of 2.6 (95 percent CI: 1.3–3.9). The gap between Malawi’s multiplier of 1.4 and Kenya’s 2.0 is explained by Kenya’s lower openness; both economies are highly slack, but Malawi is more open than Kenya.

Beyond slackness and openness, the model highlights an additional margin that can affect the size of the multiplier: sectoral reallocation. Our random shopping microfoundation produces a surprising (and informative) link between the extent of this reallocation and the multiplier. Transfers increase demand for non-agricultural goods and raise the return to capturing a larger share of the customer base, inducing households to shift labor toward non-agriculture in order to recruit customers. Under slackness, however, the total number of customers in aggregate is fixed by demand. This reallocation is thus “zero-sum” and, while individually rational, generates no additional non-agricultural output in aggregate. It does, however, drain labor from agriculture, leading to a lower output and thus a lower multiplier.

In light of this, we allow for a reduced-form sectoral reallocation friction where households can only close some exogenous portion of the gap between their current and desired sectoral labor split. Our preferred specification is fully frictional, where households cannot reallocate labor at all in response to the transfer, for three reasons. First, the frictionless case generates unrealistically fast reallocation — non-agricultural hours increase

by 14 percent the first month after the transfer and decline back to baseline after less than a year — violating intuition about the nature of agricultural work (i.e. one would not just abandon an already-planted field). Second, the experimental evidence closest to our context, [Egger et al. \(2022\)](#) in Kenya, finds no evidence of sectoral reallocation. Third, comparing the model’s predictions for the income multiplier — which range from 1.3 (frictionless) to 2.0 (full frictions) — to the experimental point estimate of 2.6 (95 percent CI: 1.3–3.9) suggests that the frictional case corresponds more closely to reality.³ Therefore, we use the results of the frictional model as our predictions for the eventual experimental results in Malawi. We also note that this suggests that the speed of sectoral reallocation is a natural target for future experiments that survey households earlier and at higher frequency.

In our final section, we use the model to ask how a fixed transfer budget should be split in order to maximize welfare or total income. We consider a situation where the total budget for transfers is the same as in the GiveDirectly experiment, but the policy designer can choose a combination of coverage (the fraction of villages receiving transfers) and intensity (the size of the transfer relative to average pre-transfer household income). 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 income multiplier). The two criteria give opposite conclusions. Consumption-equivalent welfare per dollar is maximized by covering as many TAs as possible at low intensity, because the curvature of the utility function rewards reaching more households. On the other hand, the income-maximizing transfer is large and concentrated in a single TA, as concentrated transfers help mitigate the (negative) impact of sectoral reallocation described above by minimizing the number of households who are incentivized to shift sectors.

Related Literature We provide an explanation for the large transfer multipliers along with minimal price inflation documented in field experiments in low-income settings. On the size of the multipliers, [Egger et al. \(2022\)](#) report a local transfer multiplier of about 2.6, from one-time transfers worth over 15 percent of local GDP across 653 Kenyan villages. [Banerjee, Faye, Krueger, Niehaus and Suri \(2023\)](#) find that Kenyan villages receiving a universal basic income saw an expansion in non-agricultural enterprises. For *Bolsa Família* in Brazil, [Mendes, Miyamoto, Nguyen, Pennings and Feler \(2023\)](#) estimate a relative multiplier of 2 to 3 at the state level, with little change in total employment so that

³It is important to note, however, that almost the entire range is contained in the experimental 95 percent confidence interval. This suggests that these estimates are likely underpowered for identifying the degree of reallocation frictions, a point which we elaborate on in the main text.

the multiplier reflects higher output per worker. [Gerard, Naritomi and Silva \(2026\)](#) also find that an expansion of the program led to an increase in formal employment.

On the effects of cash transfers on price inflation, the literature has been reporting little movement in average prices. But the same studies document sizable heterogeneity in price effects across markets or goods. [Egger et al. \(2022\)](#) find small price increases overall, with price increases concentrated in more isolated markets. [Cunha et al. \(2019\)](#) exploit the randomized design of a food assistance program in Mexico, where poor households received cash transfers, and report that cash raised food prices only by a negligible amount. However, they also find that the price effects are much larger in less developed villages, which they attribute to weaker ties to the outside economy and less competition among local suppliers. [Filmer, Friedman, Kandpal and Onishi \(2023\)](#) find that a household-targeted transfer in the Philippines raised the prices of less-tradable perishable foods, mostly in remote villages. For *Progres*a, a flagship conditional cash transfer program in Mexico, [Angelucci and De Giorgi \(2009\)](#) find no evidence that the transfer has induced a systematic increase in the average unit prices of basic staples. However, [Attanasio and Pastorino \(2020\)](#) show that while the transfer left average prices of staples unchanged, it led to a significant increase in quantity discounting. The price effects in [Attanasio and Pastorino \(2020\)](#) operate on the local retail margin, exactly the component our model endogenizes. However, firms in our model post a single CES price, so that margin of adjustment is unavailable in our model. Our prediction is thus about posted prices rather than about the shape of price schedules. Nevertheless, the openness in our model incorporates the heterogeneity of price effects depending on the remoteness of treated areas or tradability of goods.

Slackness, the under-utilization of productive inputs, has appeared in several literatures under different names. [Walker et al. \(2024\)](#) measure the utilization of labor and capital in a representative sample of Kenyan firms and build a model of capacity choice in which indivisible inputs generate slack in steady state. [Breza, Kaur and Shamdassani \(2021\)](#) document labor rationing in rural India, where a large hiring shock in the lean season raises employment without raising wages, and [Donovan, Lu and Schoellman \(2023\)](#) show that self-employment in developing countries is closer to unemployment than to a job, a means to sustain consumption while searching for wage work. [Herreño, Pinardon-Touati and Thie \(2026\)](#) estimate a steep Phillips curve for Indian formal manufacturing but find flatter supply curves when labor markets are slack. Slackness is not specific to developing countries. [Ramey and Zubairy \(2018\)](#) construct a measure of slackness using historical series data in the US and estimate the sizes of the government spending multi-

pliers depending on the state of the economy, although they find no evidence that multipliers are larger when the economy is slack. [Michaillat \(2026\)](#) develops a general theory of slack and provides evidence of slackness across developing and developed countries. Our model takes the concept of demand-constrained production from [Walker et al. \(2024\)](#) but microfounds it differently: firms wait for randomly arriving customers, so idle time stems from demand rather than from an indivisible input. The simpler microfoundation lets us embed slack in a dynamic model with saving and sectoral labor choice, so that the response to a transfer can be calculated on a transition path.

Openness connects the paper to the macroeconomic literature estimating the local and aggregate fiscal multipliers using regional variations ([Ramey and Zubairy, 2018](#); [Hazell, Herreno, Nakamura and Steinsson, 2022](#)). [Nakamura and Steinsson \(2014\)](#) estimate an open-economy relative multiplier from military spending across US regions, and [Chodorow-Reich \(2019\)](#) argues that such cross-sectional estimates only provide a rough lower bound on the closed-economy multiplier because part of the demand from a regional stimulus leaks to other regions through trade. Cash transfers in our setting are a clean version of such an experiment, since they are financed from outside the study area with no offsetting reduction in spending elsewhere. [Farhi and Werning \(2016\)](#) study such outside transfers to a region of a currency union and show that the transfer multiplier falls with the openness of the recipient economy, because a transfer raises local demand only to the extent that spending stays at home. This leakage is what openness captures in our model. Not only is our model consistent with the findings of the literature, it also sheds light on the interaction between openness and slackness. The model suggests that when the economy is away from slackness, openness raises the multiplier as the literature has documented, but lowers it when the economy exhibits a high degree of slackness.

This paper belongs to a literature in macroeconomic development that disciplines general equilibrium models with experimental evidence ([Buera, Kaboski and Townsend, 2023](#)). [Todd and Wolpin \(2006\)](#) use the Progresa experiment to validate a dynamic structural model of schooling. [Lagakos, Mobarak and Waugh \(2023\)](#) and [Fujimoto, Lagakos and VanVuren \(2023\)](#) combine randomized experiments on migration and schooling, respectively, with general equilibrium models to quantify aggregate effects that the experiments alone cannot measure. Other studies using this approach have investigated managerial delegation ([Akcigit, Alp and Peters, 2021](#)) and microfinance ([Buera, Kaboski and Shin, 2021](#)). In this paper, the model is estimated using the baseline data of an ongoing experiment. Although the model is not disciplined to match the endline results by construction, it provides a framework that can shed light on the program design and mechanism, along

with useful pre-registered predictions that can be tested with the endline data afterward.

Two companion papers in this issue make quantitative predictions for the same Malawi experiment with different assumptions about the underlying economy. In [Mendes, Miyamoto, Nguyen and Pennings \(2026\)](#), firms with sticky prices respond to the higher demand by producing more, generating a local multiplier of around one with little first-year inflation. In [Buera, Singh and Shin \(2026\)](#), the transfer relaxes the working-capital constraint, raising intermediate inputs and business entry. In our model, output is demand-constrained and idle capacity absorbs the transfer, so prices do not move. Our model does not have a labor market and hence makes no predictions on wages. The three different mechanisms are not mutually exclusive, and the endline data would be informative for quantifying their contributions. Higher intermediate-input purchase and business entry would point toward [Buera et al. \(2026\)](#), rising hours toward [Mendes et al. \(2026\)](#), and flat prices with output rising at unchanged hours toward ours. We consider the three papers as complements, each pointing to a different mechanism behind the aggregate effects of cash transfers.

2 Model

There is 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 .

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

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.

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

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

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

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 and Trade Costs: Agricultural goods are freely traded at a fixed price and serve as the numeraire. The fixed price reflects an implicit assumption that these goods are traded on a broader international market of which the model economy is only a small part and that any differences between equilibrium supply and demand are made up for by imports/exports on this market, rather than price changes. Thus this sector is best thought of as a sector of highly “commodified” goods that are largely fungible, transportable, and widely traded.⁴

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 $\tau_{v,w} p_{N,w}(i)$. Trade thus follows a classic Armington structure where varieties within a village aggregate with elasticity σ and varieties across villages aggregate with elasticity σ_f while trade costs enter via the budget constraint (below). In contrast with agriculture, this sector reflects local value-added in less-fungible goods and services (in practice, this turns out to largely be retail services).

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 an 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

⁴In Malawi, the labels of “agriculture” and “commodity” largely overlap, as the vast majority of household expenditure is on food, but we do note that there are a few consumption categories, such as fuel, that are highly commodified — and thus belong in this category — but are not agriculture.

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 θ of households to be entirely hand-to-mouth, choosing $a' = 0$ each period, rather than optimally smoothing consumption. Without this parameter, the model delivers MPCs much lower than those measured in cash transfer experiments and, consequently, severely underpredicts multipliers.

Non-Agricultural Production: 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.⁵ 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

⁵In 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.

with share parameter ω and converts $\frac{1}{z}$ units of this combined input into a unit of the household's unique variety. For simplicity, we define $c(z) = \frac{1}{z}$ so that $1^{1-\omega} P_{N,v}^\omega c(z)$ is the cost-minimizing per-unit cost of production for a z -type household in village v ($P_{N,v}$ is the CES ideal price index for the village's aggregate non-agricultural good). The parameter ω then governs the (cost) share of intermediate goods/services that stem from local sources and generate local income while $1 - \omega$ of each dollar spent on intermediates is spent on goods implicitly sourced from the international market (and outside of any village in the model).

Profits at a given quantity-price (q, p) combination are then given by

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

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

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

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.⁶ 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 uni-

⁶For local consumers, the timing of market visits is shaped by their own agricultural and domestic tasks. For traders, it is shaped by travel logistics.

formly 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 goods from village v so that $S_{N,v}$ is total spending on v 's non-agricultural goods.⁷

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 $\bar{q}_N^*(i)$, the household can accommodate increases in demand (i.e. increases in $S_{N,v}$) without additional labor input.⁸

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

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

⁷Closed-form expressions for these shares in terms of parameters and prices follow from the household's consumption optimization; see Appendix B.1.

⁸Note that in this stylized model, the household does not distinguish between the number of customers and the quantity demanded per customer — a doubling in the number of customers and a doubling of demand per customer both result in an expected doubling of the quantity demanded generated by being open in any given moment of the period. In reality, firms with a low number of high-demand customers may be better able to coordinate opening schedules than those with high customer volumes.

The income maximization problem of the household is then given by

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

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.

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

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

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

Idiosyncratic Shocks and Income: Finally, we allow households to experience idiosyncratic and symmetric shocks to both agricultural 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 B.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 equations 1–5), the household income maximization problem (described by equations 6–14), and the market clearing condition (described by equation 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 a single treated village that trades with the rest of the world and is small, so that any cash transfers have no impact on world prices.


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 (i.e. an increase in $\bar{q}$). The second (the “Open Economy”) maintains the zero-slack assumption but reduces trade costs to generate a substantial amount of trade.⁹ Figure 1 compares the impact of cash transfers on the treated village between the Baseline, Slack, and Open economies. We vary the size of the transfer from 0 to 50 percent of average village income and plot the impact on real output within the village (left panel) and local non-agriculture prices (right panel), providing insight into how the impact of the transfer varies with its size.

Examining the figure, we see that slackness and openness both substantially mitigate the price impacts of the transfer. In the case of slackness, non-agricultural prices stay completely flat until the transfer is large enough to overcome underlying slackness, at which point they also begin to appreciate.¹⁰ On the other hand, openness mitigates the increase approximately equally regardless of transfer size — one way of thinking about

⁹It is important to highlight that these results come from a simplified economy with parameters chosen essentially at random. A fully calibrated model is introduced in Section 4.

¹⁰Recall that 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



this result is that more competition from foreign producers limits the extent to which local prices can appreciate.


The impact of slackness and openness on output in the treated village is similarly intuitive. As discussed above, the slack economy possesses idle production capacity, even in equilibrium, that can be activated in response to transfers. As a result, output increases in response to transfers up until the point where slackness is exhausted. Here, all firms possess the same value of $\bar{q}$ and thus this occurs uniformly and suddenly at a transfer size of around 35 percent of average income. After this, further increases in the size of the transfer result in no additional output. The open economy, on the other hand, has no idle capacity and thus no increase in output. There is, however, a decline in output in the baseline economy (stemming from the fact that the increase in non-ag prices distorts the allocation of labor relative to pre-transfer prices), which is substantially mitigated in the open economy, as openness mitigates the price change.

Figure 2 displays how slackness and openness impact the consumption (left panel) and income (right panel) multipliers.¹¹ To emphasize how these multipliers vary with the size of the transfer, we plot the marginal (rather than average) multiplier so that, for example, the value displayed at 30 percent represents the additional consumption/income that would be gained from increasing the size of the transfer from 29 percent of average income to 30 percent.

Compared to the baseline economy, both slackness and openness increase the consump-

¹¹The income multiplier is computed *inclusive* of the transfer so that a multiplier of unity reflects the increase in income from the transfer with no amplification or compression.

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



tion multiplier and produce 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, greater than unity (though potentially barely) multiplier even in the presence of large transfers. The fact that openness can generate a multiplier greater than unity despite no increase in output stems from something akin to a terms-of-trade effect — under openness, the village produces more local goods than it consumes so that increases in local prices expand, rather than shrink, households' budget sets. These results highlight 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 very different ways. That slackness increases the income multiplier is unsurprising at this point; the increase in output from underutilization of production capacity manifests as household income and, for the same reasons as the consumption multiplier above, the increase falls off dramatically as the transfers grow larger. Openness, on the other hand, does not increase local production and thus has a smaller impact on the income multiplier. Its impact is still positive because, as discussed above, openness mitigates price increases and thus reduces the distortions induced by the transfer.

Impact on Imports/Exports: It is worth briefly remarking on how the transfer impacts the trade balance between model villages and the international market that implicitly underlies the assumption of fixed agricultural prices. Because transfers are one-time and

the economy eventually returns to the pre-transfer steady-state, it must be true that every dollar transferred eventually becomes a dollar spent on imports (in this case, agriculture); identical pre- and post-transfer steady states must have identical levels of savings, so the economy must dissave exactly the total amount of transfers received which, by something akin to balance of payments, necessitates an equal amount of imports. From this perspective, the key questions of slackness and openness translate to questions of how much production amplification and changes to the “terms of trade” happen along the way.

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 are 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 4.4 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 consumption 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 [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 1.45. 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.

Table 1: Externally Chosen Parameters

Parameter	Description	Value	Source/Target
<i>Normalizations</i>			
A	Agricultural TFP	1.0	Normalization
z	Non-agricultural TFP	1.45	Normalization
<i>Preferences and technology</i>			
β	Discount factor	0.95	Standard value
ρ	Coefficient of relative risk aversion	1.1	Standard value
γ	Ag returns to scale (DRS)	0.55	Chen et al. (2023)
σ_f	Across-village CES trade elasticity	4.0	Standard value
σ_a	Ag/ non-ag CES elasticity	0.5	Standard value
y	Transitory income process	See text	
<i>Data-matched structural parameters/Experimental design</i>			
ω	Share of local intermediates	0.39	Enterprise census
$[\tau]$	Iceberg trade cost	2.8	Home share of non-ag = 0.71
Transfer amount	Per-HH transfer / consumption	0.70	Chiradzulu PAP
N	Number of Traditional Authorities	10	RCT design

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 σ_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.¹²

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 ω and the trade cost matrix $[\tau]$. To compute the intermediate goods’ share, we rely on data from the baseline surveys that ask enterprises to report the source of their inputs directly.¹³ In total, village enterprises report that 39 percent of their inputs are sourced from within their own Traditional Authority (our definition of “local”).

Determining the trade costs is slightly more involved. 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 with 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

¹²We 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 these 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 7.5 km and most transport happens on foot), the implied value for σ_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.

¹³These data were not available to us directly, and we thank the experimental team, Daksh Walia in particular, for their assistance in computing these moments.

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 σ_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 et al., 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 (0.26, all local), and the share of intermediate goods sourced locally (0.39), matching this value requires that 71 percent of non-agricultural spending occur locally, which corresponds to an iceberg parameter of 2.80.

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 level at which treatment saturation is randomized in the trial. Within the trial, there are ten Traditional Authorities (TAs) assigned to three potential levels of saturation. Within a TA, villages are combined into location-based subgroups. Low-saturation TAs then have 40 percent of these subgroups treated where two-thirds of villages (i.e. all households within the village) receive transfers while one-third do not (and no transfers to untreated subgroups) so that 27 percent ($= 0.4 \times 0.67$) of households receive transfers. Similarly, medium- and high-saturation TAs have treatment shares of 40 and 53 percent, respectively. Thus we choose $N = 10$ with transfer density varying between units as in the experiment.¹⁴ As in the experiment, the size of transfers is set at 70 percent of annual income.

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

¹⁴Under symmetry, the level of N essentially acts as a normalization. Choosing $N = 100$ with the same distribution of saturations gives very similar results.

on savings. The value of the elasticity of substitution between household-level goods σ is chosen to match a profit-to-revenue ratio of about 31 percent as measured in the enterprise component of the household baseline survey.¹⁵

The share of households who are hand-to-mouth θ 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 hand-to-mouth share of 88 percent. The relative weight of agricultural consumption in utility, α , is chosen to match the overall agricultural consumption 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 in a manner consistent with systems of national accounts. Under this definition, we compute that agriculture makes up 26 percent of household consumption.¹⁶

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 baseline 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.

¹⁵It is worth noting that although the relationship between the markup and σ 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 σ as an SMM parameter rather than a directly estimated parameter.

¹⁶Matching 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 substantially 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
σ	3.25	Profit/revenue ratio	0.308	0.308	HH baseline
θ	0.88	27-month cumulative MPC	0.93	0.93	Egger et al. Table C.I
α	0.26	Ag expenditure share	0.258	0.258	HH baseline
$\bar{q}$	7.96	Capacity utilization	0.291	0.291	Enterprise census

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. Table 3 displays the main results, that is, the multipliers and price effects, from this exercise. We note again that there are many different potential multipliers. The first is the study-area-wide income multiplier — that is, the ratio of the increase in total household income across all TAs (both treated and untreated) to the total transfers.¹⁷ The second is the study-area-wide consumption multiplier, which records the same thing but for household consumption rather than income. For Malawi, our baseline model generates exactly zero inflation and a relatively modest income multiplier of 1.10. For Kenya, it generates a slightly higher multiplier of 1.33 and similar flat price effects. 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.

These results raise two questions, which we address in turn in the next two subsections. First, why does this baseline model generate such low multipliers in general — the multiplier in Kenya is substantially lower than the experimental point estimate from Egger et al. (2022) of 2.6 — and does this suggest any model extensions? Second, why does the model generate a smaller multiplier in Malawi than Kenya?

4.3 Sectoral Reallocation and Frictions

Our first question was why the model seems to underpredict the multiplier measured in Egger et al. (2022) (1.3 for the model, 2.6 in the data, though with a lower 95 percent

¹⁷Notably, this only counts income 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.

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 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 output/income, 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.¹⁸ 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 14 percent in the first month after the transfer in Kenya (25 percent in Malawi), 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.¹⁹ This takes the form of an additional parameter ϕ that repre-

¹⁸The 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.

¹⁹We 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 3: Multiplier as a Function of ϕ (Kenya)



Notes. This figure plots the 27-month horizon income multiplier as a function of the labor adjustment parameter ϕ , 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 and shading represent the point estimate and 95 percent confidence interval from [Egger et al. \(2022\)](#).

sents 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 model with full flexibility while the case of $\phi = 0$ corresponds to a model where there is no reallocation — frictions are maximal, and labor input shares are fixed at their pre-transfer levels.

Figure 3 plots the value of the income 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 1.3. 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 ϕ) increase the multiplier dramatically. As ϕ approaches zero (the case in which no reallocation occurs), the multiplier increases to 2.0 — much closer to the experimental point estimate and well within the 95 percent confidence interval.

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.

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 ϕ . 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 0.04 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 are underpowered.²⁰ The predicted one-month increase of 14 percent translates to 2.4 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 18 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 and Endline Predictions: What does this imply should be the model’s predictions for the eventual endline results in Malawi? Table 3 displays the key results (multiplier, price impacts, sectoral reallocation) in the case of full reallocation frictions ($\phi = 0$) and no reallocation frictions ($\phi = 1$) for both Malawi and Kenya. Though, for the power reasons detailed above, it is impossible to use results from Kenya to confidently distinguish between the two, the point estimate for the multiplier suggests that the frictional case is closer to reality. Thus our preferred specification is the model calibrated to match Malawi with maximal frictions added, and the predictions of this model, which are our predictions for the endline outcomes of the eventual experiment, are in bold. It is important to note, however, that because there is no data moment that directly pins down ϕ , these predictions are the result of a data-informed model adjustment rather than a rigorous model estimation procedure.

²⁰Egger 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.

Table 3: Model Predictions for the Malawi Experiment

Quantity	Malawi		Kenya	
	$\phi = 0$	$\phi = 1$	$\phi = 0$	$\phi = 1$
Income mult. (study area)	1.43	1.10	2.00	1.33
Consumption mult. (study area)	1.35	1.04	1.90	1.28
$\Delta P_{\text{non-ag}}$ (%)	0.0	0.0	0.0	0.0
$\Delta \text{CPI}_{\text{treated}}$ (%)	0.0	0.0	0.0	0.0
Maximum Δ Non-ag Hours (%)	0.0	24.9	0.0	14.4

Notes. Cumulative outcomes over the 27-month transition path under the Malawi RCT/estimation and the Kenya RCT/estimation. Multipliers are cumulative real changes per dollar transferred, summed over all villages in the study area (treated and untreated); the income multiplier counts the transfer itself as income. Maximum Δ Non-ag Hours is the largest percent deviation of the treated village’s non-agricultural labor share from its pre-transfer level at any point along the transition.

4.4 The Slackness Plateau


Our second question was why the model predicted lower multipliers in Malawi than in Kenya, despite the fact that neither experienced any change in prices. To answer this, we need to take a closer look at how slackness and openness interact.

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 “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, but $\bar{q}$ also never binds in any period after the transfer. In essence, this is the maximally slack case.

Figure 4: Income Multiplier, CPI Change, and Multiplier Gap with Varying $\bar{q}$.


(a) Income 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 of (h, ω) are $(0.71, 0.39)$; the legend reports home-bias as the share of spending that generates local income (0.56 for Malawi).

Figure 4 demonstrates this by plotting the income 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 ω 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.²¹ 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 in this plateau. The estimated value of $\bar{q}$ is 8.0, whereas the plateau begins at around 2.8. 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. 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 (see panel (b)). In contrast, more openness results in more leakage of 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.²²

Anatomy of Multipliers: To highlight the channels driving the difference between multipliers for Malawi and Kenya, the upper portion of Table 4 provides a comparison of the key relevant moments computed for each country. In both countries, the agricultural consumption share is very similar (0.26 vs 0.19). The share of households that are

²¹To 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.

²²Technically, 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*.

Table 4: Key Parameter Comparison: Malawi vs Kenya

Parameter	Desc.	Malawi	Kenya
θ	Hand-to-Mouth Share	0.88 ¹	0.89
$\bar{q}$	Max Non-ag Production	8.0	9.1 ²
α	Ag. Share	0.26	0.19
h^3	Local Non-ag Spending Share	0.71	0.87
ω	Local Intermediate Share	0.39	0.73
Malawi Mult. w/ Kenya Value			
Parameter		Experiment	Single TA
(Baseline)		1.4	1.3
θ		1.4	1.3
$\bar{q}$		1.4	1.3
α		1.5	1.3
h^3		1.4	1.4
ω		1.8	1.5

This table compares the key moments determining multipliers between Malawi and Kenya. Experiment is Malawi's actual full-coverage design (all ten TAs treated); Single TA is a counterfactual with only one TA treated, the rest as network controls. 1: Imputed from Kenya. 2: Imputed from Malawi. 3: Technically, h is a targeted moment which τ is chosen to match directly, but for clarity of exposition we report the value of h here instead.

hand-to-mouth (θ) and the parameter governing the extent of slackness are slightly different between the two economies but target identical multiplier-relevant moments; both economies target an MPC of 0.93 (measured for Kenya, imputed for Malawi) in disciplining θ and capacity utilization of 0.29 (measured for Malawi, imputed for Kenya) in disciplining $\bar{q}$, so these differences are inconsequential.

The relevant difference occurs in the openness parameters h and ω , which govern the share of household non-ag spending and the share of intermediate good spending that occur locally, respectively.²³ These are both higher in Kenya, but the largest difference occurs in the intermediate spending share; Malawi is more open mostly because its enterprises source more of their inputs from outside their local area. Notably, a large portion of this gap reflects differences in the use of hired labor (which, given the model's lack of a non-agricultural labor market, we count as an intermediate input provided by a local en-

²³Technically, h is a target moment that the parameter governing trade cost τ is chosen to match exactly via a one-to-one mapping. However, the value of h is the value that is relevant for the link between openness and the multiplier, so we treat it as a parameter here for clarity of exposition.

enterprise producing “labor services”). Malawian enterprises use essentially no hired labor while Kenyan enterprises spend almost 20 percent of revenue on wages. So this parameter difference reflects a fundamental difference in the structure of production between the two settings.²⁴

The lower portion of Table 4 demonstrates how these differences in parameters translate into differences in study-area-level income multipliers, providing a breakdown of the gap between the two countries. Unsurprisingly, θ and $\bar{q}$ make essentially no difference — these are chosen to match the same moment in both economies and end up at very similar values. The agriculture consumption share α , which governs the extent to which households consume goods with internationally fixed prices/flexible international supply directly, turns out to be relatively unimportant as well. This is not because it is theoretically unimportant — it is one of the parameters governing openness and Kenya’s lower value does indeed contribute slightly to its large multiplier — but simply because Malawi and Kenya end up having similar values.

The two parameters that differ the most between the economies are h , the share of non-agricultural (i.e. flexible price) spending that happens within a household’s own TA, and ω , the share of intermediate goods sourced from an enterprise’s own TA. Perhaps surprising is the fact that h alone has little impact on the multiplier and *almost all* the gap between Malawi and Kenya is explained by the differences in the source of intermediate goods. This is due to a quirk in the experimental design, which provides some transfers to all TAs in the study area, with some variation in saturation; h governs leakages from TAs to other TAs within the study area which roughly cancel out when all TAs are treated.²⁵ The second column of Table 4 displays multipliers for a hypothetical policy that provides transfers to a single TA (with full concentration), and here h does have a non-trivial impact. Still, the overwhelming conclusion of this exercise is that it is almost entirely differences in the sources of intermediate inputs that explain the differences in multipliers

²⁴Another important thing to note is that while the Malawian data measure intermediate input usage and purchase location directly, these questions were not asked in Kenya. Our value for Kenya comes from Egger et al. (2022), who estimate that 18 percent of each dollar spent locally “leaks” via imports using context-informed best-guesses of the import share of various firm cost categories rather than direct measurement. They note that their estimates should be considered an upper bound on the import share, which translates to our value of ω being a *lower* bound for the local share and thus a lower bound for the difference between Malawi and Kenya.

²⁵There is technically some variation due to differences in treatment concentration between TAs, but the experimental range of 40–80 percent saturation is not enough for these effects to be large.

between the two settings.²⁶

5 Implications for 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 and/or multiplier at the same cost?²⁷

To investigate this quantitatively, we set the policy design space as a two-dimensional grid over village coverage κ and per-household transfer intensity λ . The coverage grid ranges from $\kappa = 1/10$ to full coverage $\kappa = 1$. Recall that the experimental design in Malawi covers ten roughly equally sized Traditional Authorities so that our lower bound of $\kappa = 1/10$ corresponds to the treatment of a single TA (and we do not consider cases “smaller” than this). The intensity grid considers intensities from $\lambda = 0.20$ to $\lambda = 7.0$ times steady-state household income. For each cell, we solve a separate transition path of the calibrated model under the corresponding (κ, λ) scenario.

Whose Welfare? Until now, we have talked in terms of prices and multipliers; pivoting to welfare requires taking a stance on whose welfare is valued. For example, suppose transfers are made to a single TA ($\kappa = 0.1$). Households in this TA will spend some amount of the transfer within their own TA, some amount on goods from *other* TAs, and some amount on goods that ultimately come from the international market/outside the study area. Goods purchased from other TAs increase demand and (if there is slack) generate production in those areas, and even goods purchased from outside the study area must generate income for *someone* by simple accounting. We opt to take the stance that the first and second (i.e. welfare for those receiving the transfer, as well as those in the study area but not receiving the transfer) are included while the impact of purchases on international

²⁶Appendix Table A1 presents this same breakdown in a different way by showing how the multipliers in both Malawi and Kenya change when the channels induced by these parameters (i.e. slackness, openness in consumption, openness in intermediates) are completely shut down. The lesson is the same: slackness is an important contributor to the multiplier and, conditional on a fully slack economy, Malawi’s higher openness in intermediates *decreases* its multiplier relative to a world where the economy is fully closed.

²⁷Note 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 investigating how unconditional cash transfers can best be targeted.

markets is excluded. This could reflect the normative position of a policymaker who cares about benefiting a specific ultra-poor population or the positive assumption that international markets are sufficiently competitive/lacking in slack that the increase in income is offset by (say) the disutility of additional labor needed to produce the additional output.²⁸


With this decision made, we consider two criteria, both of which are reasonable measures of policy success. The first is the typical consumption-equivalent utilitarian welfare criterion, normalized by the amount of money spent on the policy to facilitate comparisons across policies. As noted above, we compute welfare over the entire population, regardless of the fraction that is treated, so that this measure includes the gains from positive spillovers within the study area. The second is the average income multiplier, also computed across the entire population and thus inclusive of spillovers. The welfare criterion has the advantage of speaking directly to household utility, but the curvature of the utility function biases it toward broad coverage even when a more concentrated transfer would generate more output. In contrast, a policymaker whose objective is maximizing output through cash transfers, may be more interested in the average multiplier, even at the cost of lower utilitarian welfare.

Results: Figure 5 plots average consumption equivalent welfare per dollar spent across the policy grid while Figure 6 plots the income multiplier. The horizontal axis is village coverage κ , the vertical axis is per-household intensity λ , and the color of each cell indicates the welfare gain/multiplier, with brighter (yellow) cells indicating higher effectiveness. The black dot marks the Malawi RCT (which covers all TAs and thus corresponds to $\kappa = 1$), and the white dashed curve represents an “iso-budget” line of (κ, λ) combinations that generate the same total cost as the experiment itself. In both figures, the white diamond represents the welfare-per-dollar-maximizing allocation while the white star represents the multiplier-maximizing allocation.

The welfare results reveal that the welfare-maximizing (per dollar) policy is one that maximizes 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 0.7 percent of welfare per dollar spent, compared to about 0.4 for the least-optimal policy. So efficiency is about 80 percent higher under the optimal policy. In contrast, the multiplier results indicate that the multiplier is maximized when transfers are the most concentrated (in our case, $\kappa = 0.1$, corresponding to a single treated TA). This effect is not strong enough to over-

²⁸Notably, if international markets themselves are slack/subject to large demand multipliers, then this additional demand generates welfare *somewhere*, potentially very far away.

Figure 5: CE Welfare Gain for (κ, λ) Allocations



Notes. Color encodes CE welfare gains per dollar spent as a function of the fraction of TAs receiving transfers (x-axis) and the amount of cash transferred per TA (y-axis). The dashed white contour traces the allocations with the same budget as the actual Malawi design.


come the curvature of the utility function, so a utilitarian planner would still disperse the budget as widely as possible (i.e. the white diamond), but it does reveal that the welfare-optimal result is purely the result of utility curvature, rather than any fundamental economics.

Why does the multiplier rise with concentration? Recall from Section 4.3 that the transfer reallocates labor from agriculture toward non-agriculture, and that this reallocation substantially depresses the multiplier. The key is that the individual gain from switching sectors (i.e. higher non-ag income) is linear (equation 8) while the cost (i.e. lower ag income) is convex (equation 10). Thus a doubling of the transfer *less* than doubles the sectoral reallocation that occurs and shrinks the (proportional) impact of the competitive externality.²⁹

When Does the Mechanism Matter? While useful, there is a sense in which the optimality results above fail to illuminate the broader importance of both slackness and openness;

²⁹Note that this implies that with maximal frictions so that reallocation is completely shut down ($\phi = 0$), the multiplier is invariant to the size of the transfer, at least on the slackness plateau.

Figure 6: Income Multiplier for (κ, λ) Allocations



Notes. Color encodes the income multiplier as a function of the fraction of TAs receiving transfers (x-axis) and the amount of cash transferred per TA (y-axis). The dashed white contour traces the allocations with the same budget as the actual Malawi design.

they are driven instead by utility curvature and competitive externalities. This is because Malawi is so far into the slackness plateau that there is essentially no risk of price inflation or exhaustion of capacity at current transfers (at least, according to the model) regardless of the size or number of villages targeted by transfers, so these second-order dynamics completely dominate.

To better understand the model's implications in a situation closer to the slackness constraint, Table 5 presents the impacts of increasingly large transfers (Panel A) and increasingly broad coverage (Panel B). The upper bound of the transfer size is intentionally chosen to be extremely large (24 times local income) in order to thoroughly exhaust slackness. Panel A shows inflation and the income multiplier for various transfer sizes to a single TA (0.1 of the population), and Panel B shows inflation and the multiplier for a fixed transfer size (corresponding to the 8.40 column of Panel A) as coverage is increased from 0.1 to 1.0.

Looking at Panel A, a key inflection point occurs when the transfer exceeds 8.40 times average income. This is the point at which the transfer induces enough additional de-

Table 5: Demonstration of Slackness and Openness Mechanisms

<i>Panel A. Slackness: one TA, rising transfer size</i>										
Transfer / TA income	0.28	1.20	2.80	4.00	6.00	7.20	8.40	9.60	12.00	24.00
CPI inflation, treated (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	3.21	8.85	23.29
Income multiplier	1.09	1.11	1.14	1.16	1.18	1.19	1.20	1.17	1.11	0.96
<i>Panel B. Openness: transfer fixed at 8.40, rising coverage</i>										
Coverage κ	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90	1.00
CPI inflation, treated (%)	0.00	1.28	3.09	5.23	7.80	10.97	15.07	23.04	41.43	62.20
Income multiplier	1.20	1.20	1.19	1.19	1.18	1.18	1.17	1.16	1.14	1.12

mand to exactly saturate all of the productive capacity.³⁰ Up to this point, there is no inflation and the multiplier increases with the size of the transfer for the same competitive-externality-related reasons discussed above. After this point, however, the transfer becomes “too large”, and inflation begins to occur. As the transfer grows, the multiplier declines and inflation accelerates. Eventually, the price increases are so severe that the multiplier falls below unity, illustrating the risk of scaling vertically under slackness.

Panel B begins with a transfer to a single TA equal to 8.4 times local income (i.e. the seventh column of Panel A). This is already enough to (barely) exhaust local slackness, but each column increases coverage by an additional TA (i.e. 10 percent of the population). Even keeping the transfer size fixed, increasing the scope of transfers leads to higher inflation and lower multipliers, illustrating the risk of scaling horizontally when relying on openness to control price increases. Interestingly, declines in the multiplier are much smaller than in Panel A — although there are risks to scaling horizontally, they appear much less important quantitatively than the risks to scaling vertically, at least in this (admittedly) artificially manufactured example. This stems partially from the fact that non-agricultural trade with other villages, the source of openness we largely focus on, is only one source of openness in our model — international trade in agricultural goods (implicit in the assumption of a fixed price) still serves to fill excess demand.

³⁰This large number partially reflects the assumption that 100 percent of production capacity is available. Adjusting the estimate of $\bar{q}$ to reflect a maximum capacity utilization of 80 percent indicates that the slackness-saturating transfer is 6.0 times local income. The exact number does not matter much for the illustrative points here.

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 τ implies that shipping one unit of a non-agricultural good across villages requires more than one unit to be produced, with the extra share melting in transit. 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, 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 v , 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, a fraction v (so $v(\tau_{h,v} - 1)$ units in total) accrues to the exporting village as additional non-ag income. Thus the pure leakage case ($v = 0$) corresponds to the baseline while full capture by the exporting village ($v = 1$) is the upper bound.

Table 6 reports the model's predictions under pure leakage and full capture. Notably, varying v 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 an income multiplier of 1.1. The second row reports the results under full capture. Here, the income multiplier rises to 1.4 with the consumption multiplier increasing similarly (the Y-C gap also rises slightly). The mechanism is straightforward. Increasing v means that a large portion of trade costs recircu-

Table 6: Sensitivity of Model Predictions to the Trade-Income Share v .

Version	$\bar{q}$	GDP mult	ΔCPI (%)	Y–C gap	C mult
Baseline ($v = 0$)	8.0	1.1	0.0	0.05	1.0
$v = 1$	8.0	1.4	0.0	0.07	1.3
$v = 1, \phi = 0$	8.0	1.9	0.0	0.10	1.8

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

lates 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 ($v = 1$) and sectoral reallocation is shut down ($\phi = 0$). Recall that shutting down sectoral reallocation increased the multiplier by 0.3 (from 1.1 to 1.4) and, from the table, we can see that full capture of costs increases it by 0.3 (from 1.1 to 1.4). Making both changes together, however, increases the multiplier by 0.8 to 1.9, more than the sum of the two changes individually. This occurs because of slight 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 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](#)).³¹ 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 48 percent of annual income, which corresponds to a rate of return of about 9 percent annually net of depreciation. 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

³¹We note here that we have performed this exercise with a wide variety of values of depreciation, and it seems to make no substantial difference.

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 income multiplier is higher, though only marginally — it increases by 0.10. The portion of households that own a machine increases in response to the transfer and peaks at 5.8 percent, compared to 3.3 percent at baseline. The reason for these muted effects is twofold. First, the model is estimated to match an observed MPC of 0.93. Conditional on matching this moment, only a small portion of the transfer is available for investment. Second, the estimated model features 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.

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 income multiplier of between 1.1 and 1.4, with our preferred specification generating a multiplier of 1.4.

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 the 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 lowers 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 1.4.

Lastly, how best to spend a fixed transfer budget 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 1.4 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 be 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.
- Angelucci, Manuela and Giacomo De Giorgi**, “Indirect effects of an aid program: how do cash transfers affect ineligibles’ consumption?,” *American economic review*, 2009, 99 (1), 486–508.
- Attanasio, Orazio and Elena Pastorino**, “Nonlinear pricing in village economies,” *Econometrica*, 2020, 88 (1), 207–263.
- 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.
- Breza, Emily, Supreet Kaur, and Yogita Shamdasani**, “Labor Rationing,” *American Economic Review*, 2021, 111 (10), 3184–3224.
- Buera, Francisco J, Joseph P Kaboski, and Robert M Townsend**, “From Micro to Macro Development,” *Journal of Economic Literature*, 2023, 61 (2), 471–503.
- , —, and **Yongseok Shin**, “The macroeconomics of microfinance,” *The Review of Economic Studies*, 2021, 88 (1), 126–161.
- Buera, Francisco, Kuldeep Singh, and Yongseok Shin**, “Predicting the Effects of a Large-Scale Cash Transfer: Financial Frictions, Occupational Choice, and Local Labor Markets,” 2026.
- Chen, Chaoran, Diego Restuccia, and Raül Santaaulàlia-Llopis**, “Land Misallocation and Productivity,” *American Economic Journal: Macroeconomics*, 2023, 15 (2), 441–465.
- Chodorow-Reich, Gabriel**, “Geographic cross-sectional fiscal spending multipliers: What have we learned?,” *American Economic Journal: Economic Policy*, 2019, 11 (2), 1–34.
- 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.

- Donovan, Kevin, Will Jianyu Lu, and Todd Schoellman**, “Labor market dynamics and development,” *The Quarterly Journal of Economics*, 2023, 138 (4), 2287–2325.
- 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.
- Filmer, Deon, Jed Friedman, Eeshani Kandpal, and Junko Onishi**, “Cash Transfers, Food Prices, and Nutrition Impacts on Ineligible Children,” *Review of Economics and Statistics*, 2023, 105 (2), 327–343.
- Fujimoto, Junichi, David Lagakos, and Mitchell VanVuren**, “Aggregate and Distributional Effects of ‘Free’ Secondary Schooling in the Developing World,” 2023. Unpublished Manuscript, Boston University.
- Gerard, François, Joana Naritomi, and Joana Silva**, “Cash transfers and the local economy: Evidence from Brazil,” *Econometrica*, 2026.
- Graff, Tilman**, “Depreciation and Growth: Evidence from Machine Repair in Uganda,” 2026. Unpublished manuscript, Harvard University.
- Hazell, Jonathon, Juan Herreno, Emi Nakamura, and Jón Steinsson**, “The Slope of the Phillips Curve: Evidence from US States,” *The Quarterly Journal of Economics*, 2022, 137 (3), 1299–1344.
- Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi**, “Two perspectives on preferences and structural transformation,” *American Economic Review*, 2013, 103 (7), 2752–2789.
- Herreño, Juan, Noémie Pinardon-Touati, and Malte Thie**, “How steep is the phillips curve in developing economies? a sufficient statistics approach and estimates for india,” Technical Report, National Bureau of Economic Research 2026.
- 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.

- Magalhães, Leandro De and Raül Santaeuilàlia-Llopis**, “The consumption, income, and wealth of the poorest: An empirical analysis of economic inequality in rural and urban Sub-Saharan Africa for macroeconomists,” *Journal of Development Economics*, 2018, 134, 350–371.
- Mendes, Arthur, Wataru Miyamoto, Thuy Lan Nguyen, and Steven Pennings**, “Cash Transfers in Developing Countries: Local Multipliers, Aggregate Effects, and External Adjustment,” 2026.
- , —, —, —, —, —, and **Leo Feler**, “The Macroeconomic Effects of Cash Transfers: Evidence from Brazil,” 2023. Federal Reserve Bank of San Fransisco Working Paper 2024-02.
- Michaillat, Pascal**, “A Theory of Economic Slack,” 2026.
- 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 Appendix Tables and Figures

Table A1: Parameter Contributions: Malawi and Kenya

Model Change		Malawi		Kenya	
		Mult.	Δ	Mult.	Δ
Baseline		1.4	-	2.0	-
No HtM	$\theta \rightarrow 0$	1.2	-0.2	1.5	-0.5
No Slackness	$\bar{q} \rightarrow \bar{q}^*$	0.9	-0.5	0.9	-1.1
No Openness (Consumption)	$h \rightarrow 1, \alpha \rightarrow 0$	1.7	0.2	2.6	0.6
No Openness (Intermediates)	$\omega \rightarrow 1$	3.3	1.9	4.3	2.3
No Frictions	$\phi \rightarrow 1$	1.1	-0.3	1.3	-0.7
All Changes		0.6	-0.8	0.6	-1.4

Notes. This table displays the impacts of various parameter changes on the experimental income multiplier in both Malawi and Kenya in order to provide a breakdown of the contribution of various channels. The value of $\bar{q}^*$ is economy-specific, chosen so that the capacity utilization in the pre-transfer steady-state is exactly 100 percent (i.e. there is no slack).

B Model Appendix

B.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:

$$\frac{p_{N,v}(i)^{1-\sigma}}{\int l_{N,v}(j) p_{N,v}(j)^{1-\sigma} dj}. \quad (\text{B1})$$

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

$$\omega_{N,v}(i) = \frac{l_{N,v}(i) p_{N,v}(i)^{1-\sigma}}{P_{N,v}^{1-\sigma}}, \quad P_{N,v} \equiv \left(\int l_{N,v}(j) p_{N,v}(j)^{1-\sigma} dj \right)^{\frac{1}{1-\sigma}}. \quad (\text{B2})$$

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

$$\pi_{N,v}^{(w)} = \left(\frac{\tau_{w,v} P_{N,v}}{P_N^{(w)}} \right)^{1-\sigma_f}, \quad P_N^{(w)} \equiv \left(\sum_{v'=1}^N (\tau_{w,v'} P_{N,v'})^{1-\sigma_f} \right)^{\frac{1}{1-\sigma_f}} \quad (\text{B3})$$

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:

$$\xi_N^{(w)} = (1 - \alpha) \left(\frac{P_N^{(w)}}{P^{(w)}} \right)^{1 - \sigma_a}, \quad P^{(w)} \equiv \left[\alpha + (1 - \alpha) (P_N^{(w)})^{1 - \sigma_a} \right]^{\frac{1}{1 - \sigma_a}}. \quad (\text{B4})$$

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

$$\eta_w \equiv \frac{E_w}{I_w} = \frac{Ra + I_w + T_w - a'}{I_w}, \quad (\text{B5})$$

giving

$$s_{N,v,w} = \eta_w \cdot \xi_N^{(w)} \cdot \pi_{N,v}^{(w)} = \eta_w \cdot (1 - \alpha) \left(\frac{P_N^{(w)}}{P^{(w)}} \right)^{1 - \sigma_a} \left(\frac{\tau_{w,v} P_{N,v}}{P_N^{(w)}} \right)^{1 - \sigma_f}. \quad (\text{B6})$$

For the θ -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 B.2. All other terms in (B6) depend only on parameters and on the equilibrium price-and-hours array $\{p_{N,v}(j), l_{N,v}(j)\}$.

Firm demand curve. Substituting (B6) 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.

B.2 Recursive Competitive Equilibrium

Each household has permanent type (v, z, χ) , 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 θ). 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, χ) .

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 B.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 (B2)–(B4). 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 (B6).

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 μ_v is invariant under g_χ and M . Transition paths under time-varying transfer profiles $\{T_v(t)\}$ are sequences of intra-period equilibria with μ_v updated each period by one-period-ahead application of g_χ and M .

C Data Appendix

C.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 Measurement Study (LSMS) data from Malawi, Uganda, and Tanzania to estimate various properties of the 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 ($0.125^{.33}$) and the variance of an annual innovation could be approximated as 1.01 ($\frac{1.326}{1+.125^{2/3}+.125^{4/3}}$).

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 ($\frac{1.01}{1-.5^2}$). The issue is that this variance is implausibly large. For example, the implied Gini coefficient (under a log-normal income distribution) is given by $2\Phi(\frac{1.347^{0.5}}{\sqrt{2}}) - 1 = 0.59$ 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 are plagued by measurement error, and the value reflects both actual variation in underlying income and 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 an 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.

C.2 Kenya Parameters and Estimation

Here we give more details about the parameters and moments for the Kenya variant of the model. Unlike for Malawi, we possess no Kenyan microdata, so the majority of these are constructed from the tables of [Egger et al. \(2022\)](#).

Many parameters (such as discount factors and relative risk aversion) are identical across the two economies. Below, we detail only the parameters that differ between the two economies. If a parameter is not listed here, it takes the same value as in the Malawi economy, detailed in the main text. The estimation procedures for the two economies are identical; the only difference is the targeted moments.

C.2.1 Fixed Parameters that Differ

Local share of non-ag spending h : We are very thankful to Dennis Egger and Nachiket Shah, who went back to the data underlying [Egger et al. \(2022\)](#) to compute that 87 percent of sample households' non-agricultural spending occurs within their same Location (the closest geographic equivalent to a Traditional Authority in Kenya). We use this number directly.

Local intermediate share ω : Appendix Table D.1 of [Egger et al. \(2022\)](#) reports that 18 percent of non-durable spending leaks as imports. In our model this number is given by $(1 - 1/\sigma)(1 - \omega)$ (because we model profit, which is local income, explicitly). Solving for ω yields a value of 0.73. We note that [Egger et al. \(2022\)](#) report that this number is an upper bound on the share of imports, making our value for ω a lower bound which *minimizes* differences between settings.

C.2.2 Targeted Moments that Differ

Asset/income ratio: Table I of [Egger et al. \(2022\)](#) reports non-land, non-house assets as 1132 and household income as 1023 for an asset/income ratio of 1.11.

Profit/revenue ratio: Table III Panel A reports profits of 157 and revenues of 494 for a ratio of 0.33.

Agricultural consumption share α : Table I of [Egger et al. \(2022\)](#) reports that 62.2 percent of spending in their sample is on food. In Malawi, we treat agricultural purchases from retailers as consumption of retail services whose underlying agricultural component is contained in input shares, but we lack the data to separate these things in Kenya. In Malawi, we see that 69 percent of food spending falls into this category. We impute this value for Kenya, which gives an agricultural share of 0.19 ($= 0.622 \times (1 - 0.69)$).