Depreciation and Growth: Evidence from Machine Repair in Uganda*

Tilman Graff
Harvard University

March 24, 2026

Job Market Paper

Abstract

Capital depreciation restricts economic growth, yet little is known about whether it is different in developing countries and why. I study the market for machine repair in Uganda using detailed field data I collected and a spatial growth model with endogenous depreciation. I first document that the cost of machinery upkeep in the sectors I study is around twice as high as in the US, throttling investment and growth. These costs are highest for small and remote firms, which I attribute to scale effects: repair is expensive if it is only infrequently demanded. To quantify the macroeconomic implications of these findings, I build a growth model in which capital breaks occasionally and the process by which it is repaired endogenizes the depreciation rate of the economy. The model proposes a microfoundation for scale effects in the repair market based on unpredictable demand and can account for 7-9% of the variation in income between Ugandan regions when calibrated using rich survey and administrative data. I simulate and discuss the effects of counterfactual development policies aimed at increasing capital investment. Finally, I discuss implications of my findings for the measurement of global capital stocks and present back-of-the-envelope estimates suggesting low-income countries might have 15% less capital than conventionally measured.

*tgraff@g.harvard.edu, tilmangraff.github.io. I thank my advisors Emily Breza, Xavier Gabaix, Gabriel Kreindler, and Ludwig Straub for invaluable guidance throughout this project. I also thank Mohamad Adhami, Livia Alfonsi, Pol Antras, Robert Barro, Vittorio Bassi, Sage Belz, Augustin Bergeron, Gabriel Chodorow-Reich, Cassandra Cole, Lindsey Currier, Antonia Delius, Bnaya Dreyfuss, Ed Glaeser, Oleg Iskhoki, Adrian Kulesza, David Lagakos, Hugo Monnery, Konstantin Poensgen, Moritz Poll, Aakaash Rao, Felix Samy Soliman, Diego Santa Maria, Emanuel Schertz, Elena Spadini, and Shreya Tandon for insightful discussions and support along the way. I especially thank the study participants for giving their time, as well as Yafesi Namakajo, Daniel Mugalu, Robert Mpiira and the entire field team for excellent collaboration during data collection. This study received ethics approval from Harvard and HAUREC-Uganda and generous financial support from the Chae Family, the Ross Garon and Anna Suh Family Foundation, and the Griffin Economics Research Fund.

1 Introduction

Firms in poor countries operate fewer machines than those in rich countries. They employ about 42 times less agricultural horsepower per acre of farmland and import 100 times fewer manufacturing machinery per capita than their highest income peers.[1] Given the important role machinery plays in increasing productivity, the lack of physical capital in developing economies is often listed as a key contributor to differences in living standards around the world (Hsieh and Klenow, 2010).

Why aren't there more machines in lower-income countries? In the absence of frictions, the textbook answer says that the *net return on capital*, or how much a firm can expect to earn from adding an additional piece of equipment to its production process, must be lower than in richer countries. There are two reasons the net return on capital can be low: the gross return on capital (often written as r) might be low, or the costs associated with having more machinery because of added repairs, maintenance, and depreciation (often written as δ) might be high. Of these two explanations, the former has received significant attention in contemporary macro and development economics. Efforts to measure r at the micro-level using randomized controlled trials have frequently found staggeringly high rates of return (see de Mel et al., 2008, among many others).[2] Macro-level studies in turn have explored ways to reconcile this apparent contradiction by focusing on costs and frictions associated with purchasing and employing capital in developing countries (see Banerjee and Duflo, 2005, for an overview).[3]

Much less is known about the other half of the net return on capital: the costs to upkeep machinery δ. In fact, existing development accounting exercises typically assume δ to be constant across the world or within capital classes, often parameterized at their conventional US values.[4] Nevertheless, a number of factors might plausibly render the cost to keep an otherwise similar machine intact different in developing countries, such as barriers to access adequate repair, poor maintenance, low wages, or broader environmental or engineering differences. What these channels add up to has direct implications for the size of δ, the net return on capital, and ultimately the nature of investment and the path of development. Yet despite its ex-ante ambiguous direction and potentially important ramifications, to date the literature has produced only limited empirical measurement or theoretical treatment of how expensive it is to keep machines running in developing countries, why that is, and what it implies for global growth and development. My paper seeks to begin closing these gaps.

In this paper, I present micro evidence and macro implications of what happens when physical capital breaks down in Uganda. I conduct a large field survey of capital-intensive microenterprises across the country, which I use to construct a new and easily interpretable measure covering the full suite of depreciation-related expenditures firms face. Far from being a constant parameter, I find that these costs are about twice as high as comparable estimates from the US and are decreasing in firm size and local economic development. I argue that capital upkeep is subject to economies of scale: It is more expensive to maintain the first unit of capital than each following one, both within firms and across local markets. I formalize this intuition using a structural model of sudden capital breakdown, providing a new microfoundation of agglomeration effects in markets shaped by unpredictable and timely

[1] See Figure A.1 in the appendix for details about the global differences in capital intensity between rich and poor countries.

[2] Studies documenting high micro-level returns on capital in developing countries include Udry and Anagol (2006); McKenzie and Woodruff (2008); Kremer et al. (2013); Balboni et al. (2022); Caunedo and Kala (2023); Bari et al. (2024).

[3] The literature investigating reasons for low effective returns to capital includes Lucas (1990); Mankiw et al. (1992); Hall and Jones (1999); Acemoglu and Zilibotti (2001); Caselli and Wilson (2004); Hsieh and Klenow (2007); Caselli and Feyrer (2007); Hsieh and Klenow (2009); Chari and Rhee (2021) among many others.

[4] All papers mentioned in the footnote above assume constant depreciation rates across countries (except Lucas, 1990, who does not discuss depreciation).

demand. Embedded in a general equilibrium spatial growth model, these forces endogenously lead the macroeconomic depreciation rate to be higher in less developed and more capital-sparse areas, reducing investment, and exacerbating productivity differences, consistent with the empirical evidence. Carefully calibrated to the Ugandan economy, the model accounts for 7–9% of income differences across Uganda. Counterfactual simulations reveal that policy interventions in the repair market have positive aggregate effects. Lastly, I argue that overlooking differential depreciation leads to potential mismeasurement in capital stocks across the world. I present a revised global capital series and show it is around one quarter more predictive of global income differences than existing development accounting estimates.

I proceed in four steps. First, I empirically document that not having access to effective capital repair is an important drag on firm performance across Uganda. The data for my study come from two representative field surveys I conducted across 1,408 firms in three regions of the country, covering urban centers and remote local village economies. The first survey focuses on firms in motorcycle transport and food processing (such as grain mills and coffee processing), two capital-intensive sectors central to Uganda's economy. In both sectors, capital is mainly repaired by externally hired mechanics as firms are too small to have their own in-house repair personnel. The second survey focuses on these mechanic firms and allows me to trace out what drives costs in the supply of repair.

Using the rich survey data, I construct a novel comprehensive measure of "replacement investment" to quantify the impact of capital breakage on firm profits. This measure captures how much money firms need to reinvest over the course of a year to keep a machine running and in good condition, divided by its market value (a concept first discussed by Feldstein and Rothschild, 1974). It includes all expenditures on repair and maintenance, the occasional need to fully replace a broken machine, as well as the revenue firms lose due to disruptions to their business while waiting for a machine to be repaired. For example, if the owner of a $2,000 grain mill spends $200 a year on capital upkeep in the form of repairs or maintenance, the resulting "replacement investment rate" (denoted by Δ) would be 10%. Similarly, if she spent no money on upkeep, but there was a one in ten chance the machine broke and became unusable over the course of the year, Δ would also be 10%. If both forces are present, Δ compounds to 20%. I argue that my measure offers several advantages over the way capital depreciation is usually quantified (the drop in machines' second-hand prices as they age), as it directly captures firms' realized upkeep costs and does not require strong assumptions on thick and frictionless capital resale markets, which might be particularly unlikely to hold in developing country contexts. It furthermore includes quantities such as expenditures on routine repair and maintenance, which are not usually counted as depreciation investment, despite their direct importance for the net return on capital.[5] Lastly, this flexible measure is not an engineering constant, but is endogenously shaped by the local availability, price and quality of repair – key components of firms' investment considerations.[6]

Second, I use my measure to uncover three new stylized facts about capital depreciation in Uganda. First, I show that firms' expenditure on capital upkeep is high: replacement investment rates between 37% (motorcycle transport) and 55% (food processing) imply firms on average incur costs equivalent to a third to half of a machine's value each year just to keep it operational, twice as high as benchmark

[5]The System of National Accounts treats routine maintenance and repairs as intermediate consumption and not investment. It acknowledges, however, that this distinction is "not clear-cut" and often difficult to enact in practice. (United Nations, 2009, 6.226ff and A4.60). Even more major repairs that *are* counted as investment will likely not be fully picked up by machinery resale price movements, despite their immediate contribution to the costs of keeping capital running.

[6]In section 3.3, I provide qualitative evidence that the availability of repair and spare parts is among the most commonly stated concerns for firms considering new machinery investments in the Ugandan sample.

numbers derived from engineers' estimates in comparable sectors in the US. The primary contributors to these high rates are large costs of repairing and maintaining machines (11% to 25% of a machine's value each year) and substantial scrap and replacement rates of existing equipment (on average every 4 to 7 years). Especially in food processing, the market for repairs further appears plagued by difficulties in getting repair quickly as firms often wait a considerable amount of time to get their machines repaired. Over the course of a year, 8% of a machine's value is lost as forgone revenue while it lays broken and idle in this sector.

The second stylized fact is that keeping capital running is subject to scale effects and thus more expensive in remote and less developed areas. Surveyed firms in the most sparsely populated locations in Northern Uganda spend almost twice as much to upkeep otherwise similar machines compared to firms in dense urban Kampala, reaching replacement investment rates of over 70%. This gradient stands in stark contrast to the common assumption of constant depreciation rates across the economy, and is driven by remote mechanics charging higher relative prices, machines breaking more often, and having to be replaced at a younger age. Additionally, mechanics in rural areas are slower at routine repair tasks and have less subject-specific knowledge. To overcome these challenges, firms frequently hire mechanics from far away to travel and provide their repair services, incurring costly delays and monetary costs. Similar benefits to density also apply across capital classes: rare machine types (such as uncommon motorcycle brands) are more difficult to get repaired quickly than frequent ones.

A final stylized fact is that smaller firms pay relatively more to keep their machines in good repair. I show suggestive evidence that this gap can partially be explained by relatively larger firms vertically integrating and hiring in-house repair personnel to take care of their machines at lower costs. This investment, however, is only economical for firms with enough capital to generate a steady flow of in-house repair demand. Indeed, firms hiring an in-house mechanic have about a 20 percentage points lower replacement investment rate, all else equal. A caveat to this finding is that only few firms in my sample are large enough to make this investment, and so differences are imprecisely estimated (9% of firms have four or more employees, 4% have in-house repair personnel). Together, these stylized facts suggest that small and remote firms, enterprise characteristics commonly associated with developing country settings, are particularly exposed to high costs of capital breakage.

To rationalize these reduced form facts, my third contribution is to present a structural model of the economy where capital breaks down from time to time. Once a machine breaks, firms rely on the repair market to get their equipment operational again. Informed by the empirical evidence, I propose a novel theory of scale effects in markets shaped by unpredictable demand, such as the market for machine repair. The key feature of the model is that repair services are assumed to both be *unpredictable* (the need for machine repair arises randomly and cannot be scheduled in advance) and *timely* (once a machine has broken down, repair needs to happen quickly). For example, a grain mill operator in a remote village does not know when her machine will break next; but once it does, every hour she waits for the mechanic to show up carries a large cost in the form of lost revenue. I show analytically how this unique demand structure gives rise to a novel form of agglomeration effects, whereby firms in central locations benefit from the demand generated by their neighboring firms. Intuitively, if demand for repairs is higher, the equilibrium rate at which the local market satisfies this demand is faster, benefiting everyone.[7] The same forces rationalize why only large firms hire an in-house mechanic: when lots

[7] An early literature debated the question of whether unpredictability of demand rationalises increasing returns to scale (Chamberlin, 1948; Whitin and Peston, 1954; Levhari and Sheshinski, 1970; Arrow et al., 1972). Other markets share the structure of demand being unpredictable but timely, most notably emergency health care. The model would predict such

of machines might break at any time inside the firm, it becomes more economical to always have a mechanic on stand-by, rendering every single breakdown less disruptive.

The model formalizes this intuition using insights from queuing theory. Demand materializes as a stochastic process, the long run mean of which is known to mechanics, but not the exact timing of the next realization. A mechanic of grain mills might know that on average, four customers walk into her store per week, but she doesn't know when exactly. This can lead to a situation in which multiple orders come in in close succession, pile up, and form a potentially weeks-long queue. Customers dislike waiting in line as they lose revenue sitting on a broken machine, and so shift their demand to places with faster service. In spatial equilibrium, mechanics in central areas optimally respond by hiring additional workers which allows them to clear orders faster. In remote areas with low overall demand, however, these workers sit idle between orders and so generate high average operating costs to mechanics, driving up repair prices. In equilibrium, the lower overall demand for repair services, the higher the price for each service.

I embed this force in a general equilibrium spatial growth model, in which the cost of repair endogenizes the effective level of depreciation and subsequently the resulting amount of investment. In the model, monopolistically competitive mechanics are able to offer cheaper prices and faster service if they operate in high-demand areas, lowering expenses necessary to keep machines running throughout the year. Capital owners take this market structure into account when deciding how much to invest in, repair, and maintain their capital. This leads to a multiplier effect whereby capital-rich regions experience a cheaper rate of upkeep, attracting even more capital, and so on. On the other hand, remote regions find it hard to grow, as initial units of machinery fail and repair is expensive to come by.

I estimate the model using my rich microdata on the market for repairs, supplemented with administrative data from various Ugandan sources. As my survey was designed with the eventual model of frictional repair markets in mind, many parameters can be calibrated non-parametrically. Detailed data on the travel behavior and implied cost functions by mechanics, self-reported frequencies of capital breakage, as well as tax data on the trade in final goods across Uganda pin down the structural and spatial parameters of the model, shaping the relationship between local economic development and endogenous depreciation. The estimated model matches a series of untargeted moments, such as the distribution of capital imports and employment shares across the country as well as the observed gradient in replacement investment rates between more and less developed regions.

Fourth, I argue that understanding capital depreciation as an economic object subject to economies of scale has important macroeconomic implications for growth and development. In particular, if poor regions face higher repair costs they invest less, making them even poorer, and so on. I use the calibrated model to quantify how much this multiplier effect contributes to income differences across Uganda.[8] I perform a canonical development accounting exercise to show that this new channel can explain between 7% and 9% of the variation of income per capita between Ugandan regions. In other words, capital upkeep being more costly in less developed areas accounts for slightly less than one tenth of the Solow residual across Uganda. As differences in these costs within Uganda mirror differences between Uganda and the US, a similar contribution might be plausible across the world.

If high depreciation rates are inhibiting growth, what tools do policymakers have to bring them down?

markets to be subject to similar forces, a point made at least as early as Arrow (1963).

[8]Uganda displays a *within*-country variation in living standards of similar magnitude as the *across*-country variation globally, a trend in line with existing evidence on sub-national income differences (Acemoglu and Dell, 2010).

I conduct three counterfactual simulations of policies targeted at increasing capital investment. First, a nationwide vocational training campaign aiming to increase mechanic productivity by rates found in well-identified labor market experiments (Alfonsi et al., 2020) would reduce replacement investment rates by about 2 percentage points (down from an economy-wide average of 51%) leading to a 2% increase in total output. A budget-neutral removal of the current 15% Ugandan tariff on spare parts would generate almost twice those gains. Lastly, an infrastructure program cutting trade and travel costs by 10% leads to modest 1.5% improvements in overall output and about a 1 percentage point decrease in aggregate depreciation, as remote firms are better able to access more mechanics.

In the last part of the paper, I discuss how not accounting for depreciation potentially being higher in poor countries might lead to bias in the way capital stocks are traditionally measured around the world. The commonly used perpetual inventory method requires an estimate of the depreciation rate, which is frequently assumed constant across space and time. I argue that if poor countries have higher depreciation rates than traditionally assumed, they also have less capital than conventionally measured. This is because more of the observed capital investment goes towards replacing depreciated old capital. Using a series of strong assumptions, I provide a back-of-the-envelope estimate on the magnitude of this bias by linearly extrapolating US-Uganda differences to the rest of the world, and recompute a revised capital series for all countries in the Penn World Table (Feenstra et al., 2015). Capital stocks are revised downwards between 10% to 30% for the poorest countries following this adjustment as their depreciation rates are particularly high. Because of this, the accounting importance of the capital-output ratio in explaining global income differences increases by a quarter (from 20% to 25%).

I contribute, first, to a long literature on the role of capital in growth and economic development, which has sometimes struggled to reconcile micro and macro-level evidence. While the Solow (1956) model assigns capital accumulation a central role in economic growth, rigorous development accounting exercises have noted that capital-output ratios explain only a modest share of global income differences (Mankiw et al., 1992; Barro and Sala-i Martin, 1992; Hall and Jones, 1999; Acemoglu and Zilibotti, 2001; Caselli and Feyrer, 2007; Feenstra et al., 2015). This stands in contrast to a series of well-identified micro-level studies, which have documented large gains from providing firms with access to capital (de Mel et al., 2008; Caunedo and Kala, 2023; Bari et al., 2024), putting renewed emphasis on the Lucas (1990) puzzle of why capital doesn't flow to the poorest countries. Importantly however, the vast majority of studies in this area assume the costs to upkeep capital to be constant across places and the process of development. By presenting novel evidence that these costs in a developing country are plausibly much higher than previously thought, my paper offers one step towards reconciling these findings. Micro-identified gross returns to capital may appear inflated when ignoring the equally high costs of capital upkeep.

Second, I contribute to the literature on capital depreciation and its measurement. I revisit an old debate on whether the rate at which capital depreciates should be considered a constant (as argued in Jorgenson, 1963) or an endogenous economic object (as advocated by Feldstein and Foot, 1971; Feldstein and Rothschild, 1974). Empirically, I find evidence for the latter view, with measured replacement investment rates varying widely across space and between firms.[9] By providing detailed measures of the cost of capital upkeep in Uganda, my paper adds to a surprisingly small literature estimating capital depre-

[9] On the theoretical debate about the right way to measure depreciation and capital stocks, see also Cramer (1958); Griliches (1962); Bitros (1976); Hulten and Wykoff (1996); Jorgenson (1996); Fraumeni (1997); Barro (2021). Lanteri (2018) and Rampini (2019) present evidence for how the standard way of measuring capital depreciation – the fall of resale prices for used capital as it ages – is complicated by resale markets being frictional.

ciation in a developing country context. Bu (2006) uses firm-level estimates of the replacement value of current capital stocks from World Bank enterprise surveys in seven developing countries and finds that these valuations decline fast over time, implying capital depreciation rates about two to eight times those in rich countries. Schündeln (2013) conducts a similar exercise among large manufacturing firms in Indonesia, which are found to have rates comparable to the United States.[10] I add to this literature first by proposing a new and direct way to measure the cost of capital breakage to firms (building on Feldstein and Rothschild, 1974) and presenting newly collected field evidence that this cost is high.[11] I then discuss implications of mismeasured depreciation rates on the way capital stocks are constructed around the world, and provide a back-of-the-envelope revision of the Penn World Table capital series (Feenstra et al., 2015), which amplifies the role of capital in accounting for global income differences.

Third, I add to the large body of work on the growth barriers faced by small enterprises in developing countries. Research in this area has emphasized credit constraints (Banerjee and Duflo, 2005; de Mel et al., 2008), market matching frictions (Startz, 2016; Vitali, 2022), high demand volatility (Pelnik, 2024), or input indivisibilities (Bassi et al., 2022; Kaboski et al., 2024; Walker et al., 2024). My findings add to this list by emphasizing that small firms face barriers to keeping their machines in good repair, leading to high rates of capital depreciation, especially in remote areas. Part of why these costs are so high is that machine breakdown frequently implies complete production standstill in small firms, linking my findings to the influential literature on production complementarities and weak links (Kremer, 1993; Jones, 2011). Theoretically, I propose a new microfoundation of scale effects based on the unpredictability of repair demand, rationalizing why these costs are particularly high for these enterprises and characterize a novel source of agglomeration effects in dense areas (Marshall, 1890; Miyauchi, 2024; Caunedo et al., 2024). Similar to Lagos (2000), Sattinger (2002), or Brancaccio et al. (2024), I harness insights from queuing theory to characterize how dense areas are better able to provide timely service. Lastly, I embed this friction in a spatial growth model building on recent advances making accumulation dynamics tractable across space by Kleinman et al. (2023) and Bilal and Rossi-Hansberg (2023). I add to this literature by endogenizing the rate at which capital breaks down as a function of geography and investigating its effects on economic growth.

In the remainder of the paper, I outline my survey data collection (Section 2), present evidence on replacement investment rates in Uganda (Section 3), and build a model of capital repair to rationalize these findings (Section 4). I then calibrate the model to the Ugandan economic geography (Section 5), and discuss implications for development (Section 6) and development accounting (Section 7).

2 Data and context

To quantify the cost of capital upkeep in a developing country context, I conducted two surveys of the market for machine repair among 1,408 total firms across Uganda. The surveys focused on two capital-intensive industries central to Uganda's economy: food processing and motorcycle transport. Firstly, the "firm survey" (N=861) was targeted at firms operating machinery in these industries, namely food processing enterprises such as grain mills, coffee processing firms and motorcycle taxi drivers.

[10] In recent work, Caunedo and Keller (2021) investigate the adjacent concept of capital obsolescence across the development spectrum, finding it to progress more slowly in poor countries.

[11] A small and old literature analyzed theoretically the question of optimal capital longevity over the development process. Most notably, Blitz (1958, 1959) argues that poor countries should invest in less durable capital and use cheap labor to maintain it, partially consistent with my findings (see also Lewis, 1955; Hirschman, 1958). Sen (1962) similarly predicts that used capital should systematically flow to poor countries where it is of more use than scrapping it in rich countries. Most machines in my study setting are purchased new.

Secondly, a "mechanic survey" (N=547) targeted enterprises specialized in repairing these machines, i.e. food processing machine mechanics and motor vehicle mechanics.

Geographical scope and sampling. The surveys were conducted in three regions across Uganda: the capital Kampala with suburbs, the rural but well-connected Luwero district a few hours north of Kampala, and remote Napak district in the far North-East of the country. Figure A.3 in the appendix presents maps of the three study sites. Chosen to represent the full distribution of Uganda's economic geography, the regions vary widely in local population density, average income, and overall economic activity, enabling me to analyze the extent to which local market thickness shapes capital repair patterns. The study regions are defined on the basis of districts, Uganda's first-level administrative division, which are further divided into parishes. A team of enumerators visited all parishes of Napak and Luwero district (37 and 93, respectively) and 109 parishes in Kampala and its suburbs.[12] In each parish, the team approached every enterprise operating a grain milling or coffee processing machine, as well as every firm specialising in the repair of these machines.[13] Enumerators furthermore went to the population weighted center of each parish and approached the two motorcycle transport businesses and two motor vehicle mechanics closest to their starting locations that met the inclusion criteria. Appendix section C.1 explains the sampling procedure in more detail.[14] Surveys were conducted between April and July of 2025, outside of the main harvest of maize and coffee.

Motorcycle transport and grain milling are prominent sectors throughout all three regions and are largely comparable in their operation. Motorcycle taxis ("boda-bodas") are overwhelmingly microenterprises of self-employed entrepreneurs, who make up around 1% of the labor force in Napak, 2% in Luwero, and 3% in Kampala, according to the 2014 Census, representing a dominant form of hailed passenger transport throughout the country and wider East Africa (Raisaro, 2024). Food processing employs between 0.1% and 0.7% of the Ugandan workforce. A common business type in this sector is grain milling, which consists of often small-scale enterprises specializing in the milling of mainly maize, rice, and cassava by servicing demand from local clients who bring their own dried grain to have it ground for a fee.[15] Lastly, coffee hulling enterprises (also called "coffee factories") are slightly larger operations which purchase dried coffee beans from local farmers and use machinery to separate the outer husk from the inner green bean, which is then sold on and eventually mostly exported. Less than 10% of surveyed food processing firms are coffee factories (none in Napak district), the rest operate in grain milling. Figure A.26 in the appendix prints photos of machines in the sample.

A few features make these industries insightful to study. For one, firms in both industries tend to be small and largely rely on external mechanics to provide repair services after machine breakdown. This is in line with global trends, whereby most mechanics in poor countries work outside firm boundaries (see appendix Figure A.2b). Second, both food processing and motorcycle transport are present and

[12] To avoid interfering with a concurrently running study in parts of Kampala's food processing industry, the team visited 73 out of Kampala's 96 parishes and supplemented the sample with 36 adjacent parishes from Kampala's suburbs which fall into surrounding Wakiso district. For safety reasons, one parish of Napak district was not visited.

[13] As food processing machine mechanics often do not operate out of a physical storefront presence, the team furthermore relied on knowledge from local food processing enterprises directing them to the mechanics in their area.

[14] My study is concerned with the depreciation and repair of productive capital, not that of private consumer durables. As such, only motorcycle riders who perform passenger transport as a business activity are eligible for the study, not riders of motorcycles for their own private commute. Furthermore, the team was instructed to deliberately over-sample motorcycles by manufacturers other than the market-leading brand Bajaj. All statistics reported in this paper appropriately re-weight to reflect the sampling procedure.

[15] Larger grain mills buying maize from local farmers, milling it, and selling it on were also eligible and appear in the sample. Only firms milling grain for human consumption were eligible, not animal feeds producers. To identify firms in grain milling, I adapted the procedure used in Bassi et al. (2022), who kindly shared their sampling materials.

fairly homogeneous business categories across the country, yet display considerable variation in local density, with Kampala having around factor 100 more workers of either category per square kilometer than the most remote districts (according to the 2014 Census, see Figure A.4). Third, the consequences of machine breakdown are an immediate and permanent loss of income: in both motorcycle transport as well as food processing, firms cannot recoup lost demand in future periods, rendering the continued operation of machinery of first order importance.

Survey design. Firm and mechanic surveys were designed to capture the market forces shaping the demand and supply of capital repair services, as well as quantifying the sources and implications of potentially high costs for capital depreciation and economic growth. On the firm side, I collect detailed, machine-level characteristics such as age, price, estimated resale value, and machine make or model, as well as measurements of maintenance investments, utilization, fuel efficiency, and a precise history of repair and replacement events. In particular, the survey elicited the frequency at which each machine experienced operation breakdowns in the last three years, as well as the causes and consequences of the most recent such incident. Firms provided information about whether machines were repaired, who repaired them, at what costs, how long it was broken and idle until it could be operational again, and how well the enterprise was able to operate in the meantime.

On the mechanic side, the survey was designed to quantify both spatial and temporal constraints in the market for repairs. First, mechanics were asked to provide a history of business trips they had undertaken in order to repair machines on their clients' premises, including information on where they traveled, how long it took them, and how much was earned. They also completed a hypothetical exercise about the price they would charge if a customer asked them to travel to various randomly drawn locations throughout the country. Second, in order to quantify temporal mismatches, the survey elicited information about the duration of a common repair task, the frequency at which new clients' orders arrived, as well as how many spare parts are currently in store. Lastly, the survey elicited motorcycle mechanics' sector-specific knowledge by administering a short subject-relevant skills questionnaire.[16]

Sector descriptives. Table A.1 in the appendix reports firm characteristics across sectors in the sample. Firms of all four types are mainly micro-enterprises, with food processing firms, motor mechanics and food mechanics having slightly over two employees on average, while motorcycle riders are almost exclusively single-person firms. Monthly revenues range between roughly $200 for motorcycle riders, $400 for mechanics of either sector, to about $1,300 for food processing firms with between half and two thirds being offset by directly measured firm costs.[17]

Firms show signs of low and uncertain demand. This seems particularly severe in food processing repair, where firms only serve an average of two customers a day. Capacity utilization in this sector, measured through a direct survey question akin to the U.S Census of Manufacturing, is slightly above 50%. This is between five and ten percentage points lower than firms in the other sectors. Firms in this sector also report large volatility in revenue: on average, they earned about four times as much in the best week of the last month, compared to its worst week.[18]

[16]The skills test was adopted with permission from Alfonsi et al. (2020), who developed a sector-specific test for motor vehicle mechanics together with Ugandan vocational training institutes and who kindly shared their survey module. Because of budget limitations, only a shortened two-question version of their original seven-question module was asked. No such questionnaire was administered for food processing mechanics.

[17]All monetary values in this paper are converted into USD from Ugandan Shillings using the exchange rate on April 1, 2025 of 1 USD $\approx$ 3614 UGX.

[18]Measures of firm utilization were adopted with permission from Walker et al. (2024), who report results from a compa-

Turning to capital, food processing firms operate on average 1.6 machines, which are mostly grain mills, grain and coffee hullers, sealers (to seal bags of flour), and electrical weighing scales. On average, these machines cost about $2,000 when first purchased and 70% are manufactured in Uganda. Motorcycle riders, by comparison, all operate exactly one machine – their motorcycle. These are overwhelmingly imported, mainly from India (who manufactures the market-leading brand Bajaj), China (Haojue), and Japan (Honda).[19] In both sectors, firms mostly own their machines and rarely rent them out, on average less than ten minutes a day.[20]

Further data sources. Throughout this paper, I use a series of additional data from various sources. I rely on administrative data from the Ugandan Revenue Authority (URA) on the universe of value-added tax receipts since 2014 to construct a representation of Uganda's firm network across districts, as well as measures of local value-added GDP and capital imports. By definition, this dataset only covers formal activity, which is likely more capital intensive and covers around half of Uganda's true GDP (Henning and Okello, 2024).[21] Furthermore, I use the 2014 version of the Ugandan Census to calibrate the topography of employment across the country, and scrape freely available data on prices of new and second-hand motorcycles on a large Ugandan online marketplace, as well as travel distances between districts using OpenStreetMaps.

3 Motivating evidence on capital depreciation in Uganda

I use the firm survey data to quantify the effects of machine breakdown on firm performance and present a series of new stylised facts about the cost of capital depreciation across Uganda's economic geography. I first present descriptive statistics on the various cost components of capital breakdown for Ugandan firms, before compiling them into a novel and systematic measure of replacement investment, akin to a depreciation rate. Patterns in replacement investment later motivate a model of endogenous depreciation and economic growth.

3.1 Descriptive evidence on capital breakdown and its components

The fact that machines are susceptible to breakdown and malfunction impacts firms through a number of different channels. They incur expenses maintaining machines before they've broken down in the first place, repairing machines once they stop working, replacing those broken beyond repair, as well as dealing with disruptions to their business while machines lay idle. As in many developing country contexts, most activities related to machine repair in Uganda are mediated through markets, as firms are too small to have a dedicated mechanic on staff. In fact, none of the motorcycle transport microenterprises and only 8% of food processing firms in my sample have a dedicated employee tasked with maintaining and repairing capital. Whether or not this market functions well thus has direct implications for firm performance and aggregate investment. The survey data is designed to capture each of

rable exercise among a representative sample of urban and rural firms in neighboring Kenya. They find an average capacity utilization of 0.62 and a best/worst week ratio of around 2.8.

[19] Figure A.5 in the appendix prints a histogram of motorcycle brands and food processing machine origin countries for all three study regions.

[20] Firm descriptives in grain milling, including the absence of a meaningful rental market, align with similar findings reported in Bassi et al. (2022)'s study of Uganda's grain milling industry (see also Bassi et al., 2023).

[21] For ethical reasons, the survey did not ask respondent enterprises' tax status. Plausibly however, the vast majority of motorcycle firms and a significant share of the other three sectors likely operate informally, especially outside the capital Kampala.

Table 1: Capital replacement investment across sectors and regions

	Motorcycle Transport			Food Processing		
	Urban	Periph.	Remote	Urban	Periph.	Remote
Panel A: Machine repair and maintenance						
Number of full machine breakdowns per year	1.0	0.6	1.6***	1.1	0.7	1.6***
Share of broken machines fixed by hired mechanic (%)	98.9	95.7	100.0	91.9	71.4	85.3
Share of total costs spent on machine repair (%)	4.1	5.5	12.0***	1.8	4.0	15.4***
Number of maintenance activities per year	19.0	18.7	22.8***	38.6	40.9	19.2**
Maintenance performed by hired mechanic (%)	85.6	77.8	85.9	54.4	32.7	26.7***
Panel B: Repair market costs						
Annual days of machine idle time if ever broken	7.0	9.0	9.3	11.4	12.4	13.5
Lost revenue during idle time (%)	98.5	97.6	94.3***	83.2	92.2	90.5*
Revenue lost to broken idle time if ever broken (%)	2.2	3.6	3.2	6.5	5.3	3.9
Mechanic travel distance (KM)	7.0	17.2	15.7***	12.9	32.3	19.5*
Had negative experience with a mechanic (%)	59.5	55.5	34.6***	42.4	26.6	29.9*
Panel C: Machine age and replacement						
Machine age in years	3.8	3.5	3.6	5.9	6.9	3.9***
Expected years until machine replacement	3.6	3.4	3.9	7.7	9.4	6.8
Fuel efficiency loss since purchase (%)	12.5	11.0	9.7	2.9	3.0	9.0***
Hours in use yesterday	10.2	9.6	6.9***	7.2	5.7	4.7***
N	274	192	78	193	208	90

Notes: Descriptive statistics of the various components of replacement investment by sector and region. "Urban" corresponds to Kampala and its suburbs, "Periph." to rural but accessible Luwero district, and "Remote" to Napak district. Stars on the remote column indicate statistically significant differences to the urban column as $p < 0.1$ (*), $p < 0.05$ (**), and $p < 0.01$ (***). "Share of broken machine fixed by hired mechanic" refers to the fraction of times a machine breaks and is repaired by someone external to the firm (as opposed to from within the firm or not repaired at all). "Lost revenue during idle time" refers to the loss of usual daily revenue (in %) each day a machine is broken and waiting to be fixed. "Revenue lost to idle time if ever broken" corresponds to the annual revenue a firm loses to all breakages and delays combined, conditional on there being at least one breakage per year. "Mechanic travel distance" refers to the road travel distance from the hired mechanic's home location to the firm's premises. "Fuel efficiency loss since purchase" refers to the relative amount of additional fuel needed to perform a usual day of operations between purchase date and now. "N" is the number of machine-observations, some firms have multiple machines.

these channels, many of which have not been quantified together in either a developing or developed country context.

Repair and maintenance. Table 1 presents summary statistics on capital repair, maintenance, and replacement. It splits the sample along sectors and districts, with columns labeled "urban" referring to Kampala and its suburbs, "periphery" referring to Luwero district, and "remote" referring to Napak district. Stars on the leftmost columns indicate statistically different values between Napak and Kampala. Panel A reports indicators on the repair and maintenance of machines. Firms in the sample experience a total breakdown of each of their machines (defined as a mechanical issue large enough for the machine to not be operational anymore) between 0.6 and 1.6 times a year. Across both sectors, the frequency of breakdown is largest in the most remote areas. Once a machine has broken down, the overwhelming majority of motorcycle riders and a clear majority of food processing firms rely on externally hired mechanics to render their machines operational again. The expenses paid for repair services is substantial, accounting for between 4% and 15% of total firm costs per machine, with the largest shares observed in more remote areas.

Firms conduct preventative maintenance to avoid machine breakdown. I classify as maintenance any activities used to prolong the lifespan of a machine while it is still operational, such as changing oil of a motorcycle or replacing the belt of a grain mill. Such activities are performed on average between once

and three times a month. In motorcycle transport, such maintenance is again mostly conducted by hired mechanics, and firms pay about $7 each time. In food processing, maintenance activities cost about $23. External mechanics only perform about 30-50% of them, with employees of the firm accounting for the rest.

Other repair market costs. Time is of the essence when repairing machines after they have stopped working. Since most firms rely on the market to provide timely repair, delay in service delivery presents a significant cost to firms. If machines break, they are on average broken and idle for between one and two weeks per year (Panel B of Table 1). There is a gradient with remoteness in both sectors, with firms in Napak waiting about two days longer on average (although imprecisely estimated). Because firms are small, they are mostly unable to continue operating during this time, losing between 83% and over 98% of their usual revenue. Adding up all breakages and delays, firms lose between 2.3% and 6.5% of total revenue each year conditional on experiencing at least one breakdown.[22]

One source for this temporal mismatch are spatial frictions, as broken machines are immobile and not every town or village has an adequate mechanic for every task. Indeed, outside of Kampala, the average machine is fixed by a mechanic who has traveled between 15 and 35 kilometers to the site of the machine. One potential reason why firms do not always rely on their nearest mechanic are relational concerns: 25-60% of firms have had a negative experience with a mechanic in the past (such as faulty installation, theft, or long delays). Partly because of this, mechanic travel is a common phenomenon across the country, leading to potentially long delays in service provision. My survey collected detailed trip-level data on the travel behavior of mechanics, summarized in Table A.2 and Figure A.7 in the appendix. These illustrate how long travel is common among mechanics, especially in food processing. Often a trip consists of a single customer, who pays a premium to be served from far away, with around a third of total repair costs accruing to mechanic travel expenses in these cases.

Machine age and replacement. Firms replace machines that would be too expensive to repair or have otherwise aged past their useful life within the firm. High costs within the repair market can manifest as high replacement rates, as firms find it too expensive to access qualified repair services and rather opt for a new machine altogether. Panel C presents evidence on the replacement behavior of capital in my sample. Machines are relatively young: the average motorcycle (food processing machine) is below four (six) years old. Especially for food processing there is again a clear spatial gradient, with machines in the remote Napak district being almost two years younger on average. Since actual machine replacement is a rare event, I asked firms about the expected remaining service life of their machines. Firms in Napak expect to replace their machines much sooner: among food processing firms, the expected time until replacement is below seven years, one to two years sooner than in the other regions, though no such gradient is observed for motorcycles.[23] 60% of firms expect to be able to sell their machine once it is replaced in the future, on average at a price 60% below its current value.

[22]The share of total revenue lost to machine breakdown is higher for firms with more complex production processes, as each new machine introduces a new possibility for failure and operational standstill (as proposed by Kremer, 1993). As such, food processing firms, especially those in urban areas where firms tend to be larger, lose relatively more of their yearly revenue to idle time.

[23]Two further potential reasons for the young age of machinery in my sample are the comparatively low age of firms (motorcycle firms have on average been in operation for 4.5 years, food processing firms for 6.5 years) and restrictions on the import of used machinery (Uganda has in place an outright ban on the import of vehicles older than 15 years). 80% of machines in my sample were purchased new, with used machines on average being 2 years old at time of purchase. High replacement rates are, furthermore, also found among firms that have replaced a machine in the past: those machines were replaced while being between 3 and 7 years of age.

Finally, machines might incur costs related to aging even before they are replaced. As a proxy for input quality decay, firms reported the amount of fuel or electricity needed to run the machine on a regular day, both now and when it was first purchased. On average, motorcycles now need about 10% more fuel to operate on a typical day, food processing machines only about 3% more.

The last row of Table 1 prints the number of hours a machine was in active use on the last day the enterprise was open. Consistent with recent work on spatial gradients in capacity utilization, machines in remote places are operating significantly fewer hours per day than those in central places, with the average Kampala machine running almost 3 hours longer than a machine in Napak (Walker et al., 2024). The findings that remote machines are breaking more often, take longer to repair, and are expected to be retired earlier, are hence not driven by enterprises running these machines past their breaking point – if anything, the opposite appears to be true.

While these differences plausibly capture the notion that accessing adequate repair is difficult in remote regions, they might also be driven by firms in less developed areas operating equipment of worse quality. However, note that enterprises in motorcycle transport drive largely the same, mass-produced vehicle makes and models across Uganda (see appendix Figure A.5). Furthermore, Table A.3 in the appendix investigates spatial gradients of two plausible correlates of inherent machine quality: purchase price and the frequency of breakdown in the first year after purchase. While machines operated in Kampala are generally more expensive than machines in the less developed Napak (column 1), this gap closes almost entirely once controlling for *where* a machine was purchased (column 2), and hence appears more reflective of local price levels than inherent quality differences. Similarly, while Table 1 reports that machines in less developed parts of Uganda experience more frequent breakdown, the reverse is true in their first year of use (where machines in Kampala experience failure almost one full time more often than Napak machines, column 3 of Table A.3). Together, this descriptive evidence guides attention to the market for repairs as a key mediator for driving the costs of capital upkeep.

3.2 A new unified measure of replacement investment.

What do these various indicators sum up to? In order to project the evidence presented so far onto a common and easily interpretable scale, I next introduce a new measure of machine-level replacement investment. In particular, I define

$$\Delta = \left[\underbrace{(1-\mathbb{1}_R)X}_{\substack{\text{repair + maintain} \\ \text{if not replaced}}} + \underbrace{(1-\mathbb{1}_R)W}_{\substack{\text{lost revenue} \\ \text{during repair}}} + \underbrace{(1-\mathbb{1}_R)\delta^I}_{\substack{\text{input decay} \\ \text{if not replaced}}} + \underbrace{(P_K - P_{K,S})\mathbb{1}_R}_{\substack{\text{discard or sell} \\ \text{and replace}}} \right] \times \underbrace{\frac{1}{P_K}}_{\substack{\text{normalized by} \\ \text{current value}}} \times \underbrace{\frac{u_{avg}}{u}}_{\substack{\text{normalized by} \\ \text{utilization}}} \tag{1}$$

as the *Replacement Investment Rate* Δ for any given machine. This measure can be interpreted as the share of the value of a machine that needs to be reinvested over the course of twelve months in order to still have a comparable machine at the end of the year (Feldstein and Foot, 1971).[24] In equation (1) $\mathbb{1}_R$ is an indicator function that is $\mathbb{1}_R = 1$ if a machine is replaced rather than repaired, X is the expenditure

[24]Feldstein and Foot (1971) argue that replacement investment is the economically most relevant notion of depreciation expenses. For the United States, they construct a measure of what fraction of the capital stock is replaced every year, similar to the last component ("discard or sell") of my Δ-measure. They do not include the other components, though Feldstein and Rothschild (1974) in follow-up work argue that all investments used to keep the capital stock constant should be included, although they do not explore this empirically.

accruing to repair and maintainance over a year, W is the revenue forgone while waiting for adequate repair service delivery, δ^I is the cost associated with purchasing more inputs as the machine declines in efficiency, P_K is the machine's current market value, $P_{K,S}$ is the price attainable on the market when reselling the machine, u is the machine's total utilization per period, and u_{avg} is the average utilization of this machine type in the economy. Δ is hence the sum of expenditures on repair, maintenance, and replacement of machines, input decay and forgone revenue accruing to broken machines, all normalized by the current market value of the machine and scaled to a fixed level of utilization.

To illustrate, assume a motorcycle of current market value $P_K = \$1000$ has an annual probability of $\mathbb{1}_{\mathbb{R}} = 0.2$ to experience a breakdown so severe it has lost all resale value ($P_{S,K} = 0$) and needs to be replaced. In the absence of any further repair or maintenance costs, this scenario implies $\Delta = [0.2P_K] / P_K = 0.2$. If, however, with reverse probability of 80% the machine has an additional lesser malfunction which can be repaired for $X = \$100$, the measure rises to $\Delta = [0.2P_K + 0.8 \times X] / P_K = 0.28$. Additional maintenance costs or lost revenue while a broken machine sits idle add to Δ in an equivalent way, as does accounting for machines that still fetch some price on the resale market when being replaced, or machines losing some of their input productivity due to regular decay. Appendix section C.2 explains each of these components and their measurement in more detail.

Discussion. Four remarks about this measure are in order. Firstly, Δ does not generally coincide with the traditional measurement of the depreciation rate δ, namely the decline in market prices of a given machine across time. Instead, it is the relative share of the total capital stock that needs to be reinvested just to keep its current productive capacity constant for one year. The traditional conception misses that (a) some machines are never sold and so resale prices are not immediately relevant for the burden repair puts on firms, (b) maintenance and repair ceteris-paribus raise the price of a machine, thereby decreasing measured depreciation rates, while still constituting a real expense for firms, (c) repair market frictions such as long service delivery times are equally a real expense for firms without showing up in price data, and lastly (d) resale market price movements can reflect things unrelated to capital repair, such as market trends or obsolescence, which are only partially informative for the productive capacity absorbed by reinstating broken machines (Hulten and Wykoff, 1996).

Second, note that this measure is the relevant one for a Solow (1956)-style capital accumulation model. In such a model environment, the appropriate notion of capital depreciation is the amount of resources that need to be invested over the course of a period to keep the capital stock constant. Intuitively, regular expenditures on repair and maintenance should be part of this quantity, as they represent economic output that is neither consumed, nor used to grow the capital stock beyond what it already was.[25]

Third, many components of Δ in equation (1) are driven by the functioning of the repair market or represent otherwise endogenous choices of the firm (such as whether to replace a machine $\mathbb{1}_R$, or how much to maintain and repair at which price X). One would hence expect this rate to vary across firms, places, and likely the process of economic development. For example, a thick repair market (with low wait times W) should decrease Δ, while high wages paid to mechanics (and hence high values for X) will increase it. My measure hence connects to an older debate on the nature and exogeneity of depreciation between Jorgenson (1963), who argues that a constant rate for δ is realistic and appropriate,

[25] Karabarbounis and Neiman (2014) define capital depreciation as the only part of gross production that is never consumed. As such, there is a tension between this notion and the National Accounts (United Nations, 2009), which classify regular repairs and maintenance of capital as intermediate input consumption, but concede that this definition is not clear-cut and often hard to implement in practice.

and Feldstein and Rothschild (1974), who present many examples in which depreciation is shaped by firms' economic decision making.

Lastly, Δ can be rescaled and reported for any given time period, such as per year or per 100 hours of use (for instance to construct a machine's "repair factor" as is common in agricultural economics, see Kay et al., 2016). To that end, my survey captures measurements on the number of hours a machine was in use, which I use to normalize the replacement investment rates across machines into one comparable indicator. In practice, for each machine type (motorcycles of different brands, grain mills, scales, etc.) I divide by the number of hours a machine was used in the past year and rescale the result by the average number for each machine type (motorcycles, grain mills, etc.) across the sample. The result is a machine-level indicator capturing how much of its value needs to be reinvested every year if it was utilized at the economy-wide average level.

To compute Δ, one needs an accurate measure of a machine's current market value, as well as its expected resale value were it to be sold. This presents a challenge in the presence of thin second-hand capital markets or machines a firm does not intend to sell. To make progress, I rely on the standard method of asking firms to estimate the current market value and the future resale value of their machines.[26] For motorcycles, where there is an active resale market, I validate firms' estimates with data I scraped from a large Ugandan online motorcycle marketplace and confirm that firms generally do a decent job of estimating the price they could get for their machine (see appendix C.3 for details).

3.3 Stylised facts on replacement investment

With a unified measure of the cost of machine upkeep in hand, I next present four new stylized facts about capital replacement investment in Uganda.

Fact #1: Δ is high. Table 2 computes Δ using equation (1) for both sectors in the survey. It splits the replacement investment rate into its four components (repair and maintenance, repair market costs, replacement, and input efficiency decay) before reporting the aggregate measure. Firms reinvest a substantial share of their capital stock each year: 37% in motorcycle transport and 55% in food processing. As a reminder, the way to interpret this number is that a food processing firm with a grain mill worth \$2,000 on average incurs expenses of about $\$2000 \times 0.55 \approx \$1{,}100$ a year to still have a machine of identical productive capacity a year later. Similarly, a motorcycle enterprise using a motorbike worth \$1,000 on average spends \$370 throughout the year.

Repair and maintenance expenditures make up a substantial share of total replacement investment in both sectors (11% and 25% of the machine's value, respectively). The majority of these expenses come from minor repairs or maintenance costs, or any action performed while the machine is still able to operate. Non-monetary costs of machine breakdown in the form of long wait times for repair services constitute a smaller, yet nevertheless meaningful component of replacement investment. Especially in food processing, a firm loses about 8% of a machine's value in revenue each year while it waits for it to be made operational again, compared to 2% in motorcycle transport.[27]

[26] Asking firms to estimate the market or resale value of their capital stock is common practice in e.g. the World Bank Enterprise Surveys (see World Bank, 2019).

[27] To reconcile this number with Panel B of Table 1, note that the average firm loses about 5.3% of its revenue to idle time if any machine ever breaks. Adding in firms who experience no breakdown and dividing by the number of machines, we get that each machine on average leads to 1.5% of revenue loss. Roughly 1/3 of revenue is spent on intermediate inputs, which presumably are not costs accruing during idle time, implying a pure annual cost equivalent to 1% of revenue. Food processing firms make around \$15,600 in yearly revenue, or around $8\times$ each machine's market value, so that the revenue loss divided by

Table 2: Replacement Investment Rate Δ across industries

	Motorcycle Transport			Food Processing		
	Uganda	*USA*	δ_{BEA}	Uganda	*USA*	δ_{BEA}
Repair and Maintenance	0.11	*0.09*		0.25	*0.18*	
	(0.01)			(0.03)		
= annual cost of minor repairs	0.09			0.20		
+ annual number of major repairs	0.66			1.08	*1.02*	
× cost each time	0.03			0.06		
Forgone Revenue during Idle Time	0.02	–		0.08	*<0.01*	
	(0.01)			(0.01)		
= days lost to idle time	5.22			8.06	*0.34*	
× average daily lost revenue during downtime	0.01			0.05		
Replacement	0.20	*0.06*		0.20	*0.07*	
	(0.01)			(0.03)		
= annual probability of replacement	0.39			0.33	*0.11*	
× (1-average scrap value)	0.53			0.61	*0.62*	
Input Decay	0.03	*≈0.0*		0.02	*≈0.0*	
	(0.01)			(0.01)		
= additional fuel required	0.03			0.01		
× average costs	1.05			3.92		
Replacement Investment Rate Δ	**0.37**	*>0.15*	*0.17*	**0.55**	*>0.26*	*0.18*
	(0.02)			(0.06)		

Notes: Sample components of the replacement investment rate Δ, across the two capital operating businesses. As in equation (1), every component is normalized by machine values and utilization. Bootstrapped standard errors using 1000 draws with replacement at the parish-level in parentheses. See appendix section C.2 for a detailed breakdown of how each of the components is constructed and measured. Corresponding estimates for US machinery printed in gray, these are rough analogues based on available engineers' estimates and other data sources, see appendix section C.4 for details. Official depreciation rates used by the Bureau of Economics Analysis and underlying much cross-country work in gray denoted δ_{BEA}, going back to Fraumeni (1997). For motorcycle transport firms, the four components appear to not exactly add up to the full Δ rate due to rounding differences. In the second panel, downtime days and average daily forgone revenue don't multiply into average annual forgone revenue as firms with higher costs to downtime experience less of it.

Replacing machines is the other large component of Δ, constituting slightly between a third and half of total replacement investment in each sector, with firms expecting to replace their machines quickly and only be able to resell their current machines at steep discounts. Lastly, input decay contributes less to overall replacement investment, at least in the dimensions measured: the cost associated with having to purchase more fuel or electricity as the machine becomes less productive constitutes around 2-3% of a machine's value each year.

How do these numbers compare to machines in similar sectors in other countries? In the absence of comparably detailed data on capital breakdowns in the US, I compile engineers' estimates and repair expenses from a variety of sources to construct Δ for US motorcycles and farm machinery. Table 2 prints these in gray, appendix section C.4 provides further details on the construction of each estimate. For food processing, I rely on the American Society of Agricultural and Biological Engineers who publishes estimates on the longevity, breakdown frequency, and repair costs of farm machinery in the US (ASABE Standards, 2011). Their estimates do not specifically focus on grain mills or coffee hullers, but instead cover a broad range of farm machinery including tractors, threshers, and sprayers, of which I take cross-machine averages. Since there are virtually no motorcycle taxis in the US, I rely on data by the

machine value is around 8%. A similar calculation holds for the other sector.

EPA and BEA which report replacement patterns and annual repair expenditures for privately owned motorcycles.[28,29] With these differences in mind, comparing estimates for Ugandan and US machinery can nevertheless offer a coarse insight into the relative magnitudes of replacement investment rates between both countries.

For both sectors, Δ in Uganda is about twice as high as in the US. Notably, this is not driven by monetary costs spent on repair, which are relatively on par between both countries, as are the frequencies of requiring a major repair (about once a year). Instead, the differences are primarily driven by lower forgone revenue during idle time and less frequent machine replacement in the US. Farm machinery in the US is on average broken for only a few hours each year, resulting in revenue losses of less than 1% of a machine's value, or eight times less than in Uganda.[30] Similarly, the annual probability of replacement is about three times lower for US machinery in either sector.

Fact #2: Δ is higher for remote firms. Next, I document that replacement investment rates are higher in remote areas. Figure 1a prints Δ pooling over both sectors for all three geographic regions in which surveys were conducted. While urban Kampala and peripheral Luwero district display similar rates above 40%, they increase to over 75% in more remote areas (leftmost point estimates for each region). This holds also when controlling for sector, the age of a machine, or the exact machine model (motorcycle brand or food processing machine type, rightmost set of point estimates). This implies that it is almost twice as expensive to operate an otherwise identical machine in the most remote and least developed parts of the country, a strong departure from the traditional assumption of constant depreciation rates. Figure A.8 in the appendix prints the distribution of all four components of Δ across space and reveals that the premium paid by the most remote regions is mainly driven by the two main contributors to replacement investment: repair and maintenance costs, and machine replacement, which are both roughly twice as large in the remote district. The same pattern is obtained when zooming in and classifying firms by the market access centrality of their parish, i.e. comparing firms in the most urban parishes of Uganda to those in the most remote ones (see appendix Figure A.10a and appendix section C.5 for details).

Two plausible contributors to this spatial gradient are the value of the machine (p_K) and the level of utilization (u), which enter in the denominator of Δ in equation (1) and are lower in rural areas. Figure A.9 in the appendix replicates Figure 1a, after adding additional controls for these factors. Importantly, the gradient is weakened, yet robust when holding these constituent components of the replacement investment rate fixed, implying that the finding that it is more expensive to keep machines running in remote areas is not just a reflection of machines there being cheaper or used less intensely.

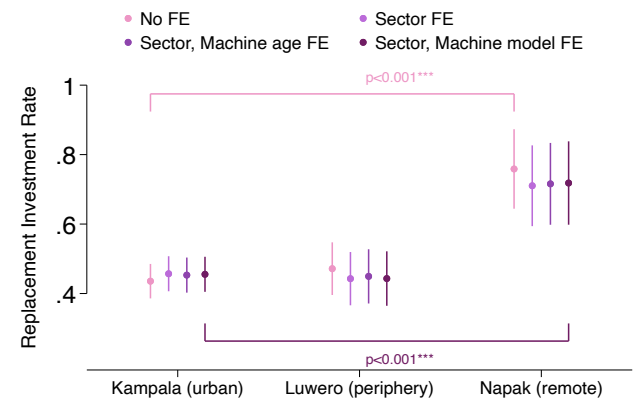
Fact #3: Δ is higher for small firms. Figure 1b prints a binscatter relationship between the replacement investment rate Δ (residualized by machine type and geographical region) and log monthly revenue of the firm owning it. Firms with the lowest revenue have rates of close to $\Delta = 60\%$, while the largest firms display about half of that. Figure A.10 in the appendix shows that same pattern holds for alternative definitions of size and splitting the sample by capital or labor inputs: firms with any paid employees or

[28] Wherever possible, I rescale US estimates to the average age and utilization rates observed in Uganda to ensure that differences along these categories do not drive the different headline numbers.

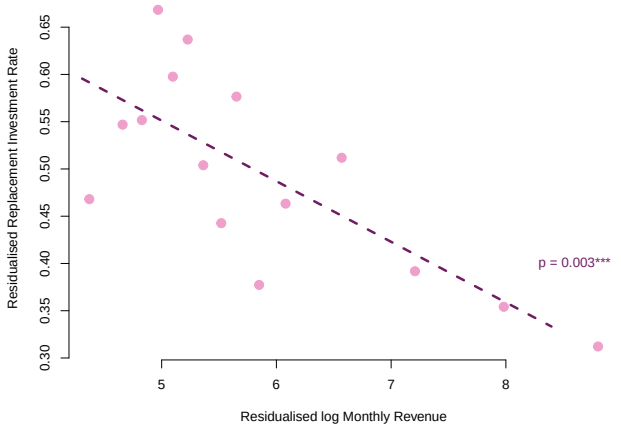
[29] Neither the EPA nor the ASABE include any evidence of input decay in motorcycles or farm machinery. In fact, the EPA refers to the conception of vehicles losing fuel efficiency as they age as a "myth" (link).

[30] Thomas and Weiss (2020) estimate the value of revenue lost to unexpected downtime in US manufacturing and report a value equivalent to about 1% of annual revenue. This compares to about 2% in Ugandan motorcycle transport and almost 5% in food processing (see Panel B of Table 1).

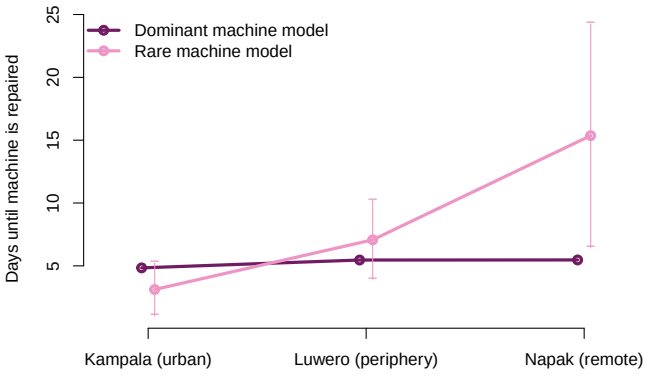
Figure 1: Stylized facts about replacement investment in Uganda



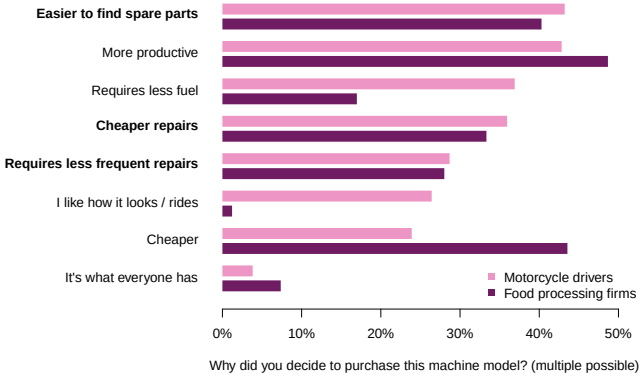
(a) Remote firms require more replacement investment



(b) Small firms require more replacement investment



(c) Rare machines take longer to repair



(d) Repairs are primary investment concern

Notes: Panel a reports the replacement investment rate Δ for firms in the three survey regions, first without any controls (leftmost estimates, light pink) before successively adding controls. As explained in the main text, Δ measures the total expenditure needed to keep a machine in constant condition throughout the year, divided by the machine market value. "Sector FE" adds fixed effects for the two sectors (motorcycle transport and food processing), "Machine age FE" further controls for the age of a machine (in 3-year buckets), "Machine model FE" adds controls for the exact machine model (motorcycle make or food processing machine type, such as grain mill or weighing scale). 95% confidence intervals printed around estimates. Brackets report p-values of a test of coefficient equality in the least and most saturated specification. Panel b is a binscatter showing that firms with higher log monthly revenue have lower Δ, controlling for geographic region and machine type). Panel c presents the average wait time to repair a broken machine by region and by whether it is the locally most common machine make (motorcycle brand or food processing machine production location). Bajaj is the most common motorcycle make everywhere and Ugandan machines are the most common food processing machine origin everywhere but Napak (where it is Chinese). Figure A.5 presents the full distribution in all locations. The gradient in delay until repair services are accessed with distance is driven by rare machine models (in pink). Table A.4 in the appendix prints these results in regression table form. Panel d reports firms' self-reported reason for why they chose their exact machine model. Multiple answers are possible. Answers relating to cost or ease of repairs boldened.

operating more than one machine have a lower replacement investment rate than firms that are owner-operated (A.10b) or operate a single machine (A.10c), respectively.

Small firms are of course different from larger ones along a number of dimensions. One important distinction is that the smallest firms are less likely to have the relevant in-house repair know-how and thus are more likely to rely on the external repair market every time a machine breaks. Indeed, firms with more than one machine (or at least one employee) are more than three times (more than 15 times) more likely to circumvent the repair market and deal with capital breakdown internally. In line with this, appendix Figure A.10d presents suggestive (and noisy) evidence that the few firms in the sample with the capability of repairing machines in-house have around 12 percentage points lower replacement invest-

ment rates than their peers without this in-house capacity. Importantly, this evidence accounts for the wages firms pay their in-house repair staff, and are hence consistent with productivity improvements in large firms.

Fact #4: Repair markets exhibit scale effects. The previous three stylized facts hint at the repair market being a key mediator to high replacement investment rates for firms. I next find evidence that repair costs are particularly severe for uncommon machinery types in more remote areas. Figure 1c prints the time it takes for a machine to be repaired after it has broken down, including time spent looking for a mechanic and waiting for the mechanic to travel to the machine. The figure then splits this delay time by machines that are of the most common brand locally to machines of all remaining brands. In particular, it compares motorcycles of the market-leading make Bajaj to all other makes, as well as food processing machines with the most frequent country of origin (China in Napak and Uganda in the other two regions) to all other ones.[31] I find that the clear remoteness gradient in the time it takes to get machines repaired is almost entirely driven by rare machine models: Bajaj motorcycles and machinery of the most common countries of origin are repaired at similar speeds throughout Uganda. Rare machine models, on the other hand, take much longer to be repaired the more remote a location is, going from three days in Kampala to almost 15 days in Napak, on average.[32,33]

Together, the evidence presented points to a repair market that is subject to scale effects. Capital in remote regions and of rare models is more cumbersome and expensive to maintain than in central areas with a busy and diverse repair market. Firms that are large enough are able to overcome these issues by circumventing the market entirely and relying on in-house repair personnel. Figure 1d presents qualitative evidence that these dynamics are also top of mind for firms' investment decisions. Asked for the main reasons why they decided to purchase their particular machine model, issues related to repairing the machine once it breaks are among the top responses by firms in either sector, with the ease of finding spare parts the most common answer among motorcycle riders (third-most common for food processing firms, behind productivity and purchase price). Whether repairs are cheap or less frequently required also rank prominently for both sectors.

Mechanic behavior Why is it so much more costly to keep machines in good condition in remote areas and when relying on the repair market? In Figure 2, I use the mechanic survey to document that repair enterprises in less developed areas are slower and less skilled than those in central areas. To maximize power, I group mechanics by the market access of their home parish and compare those in the most central to those in the most remote parishes.[34] Figure 2a reports the average time it takes mechanics

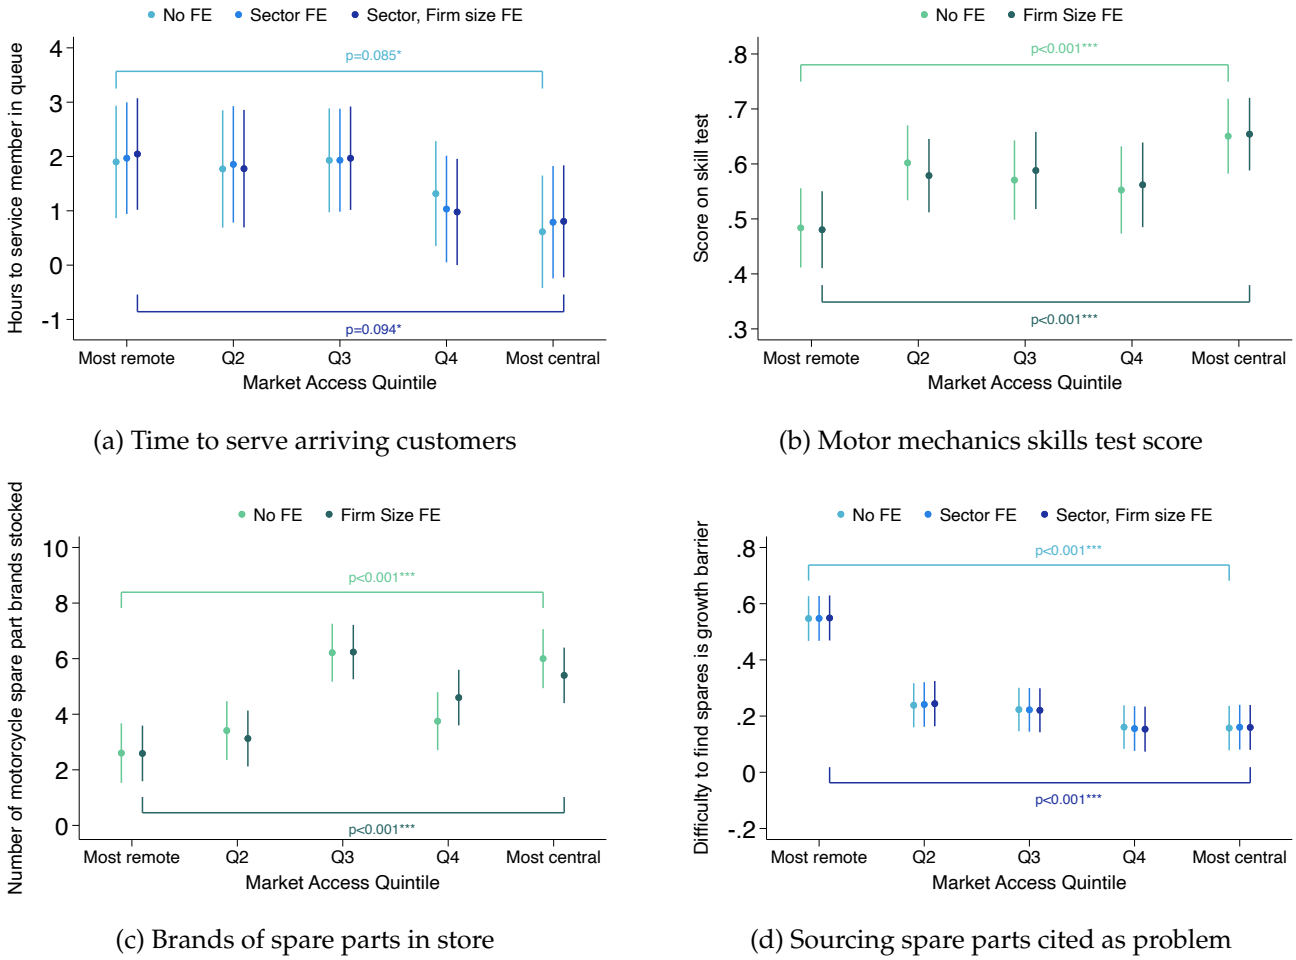
[31] Figure A.5 in the appendix prints the full distribution of motorcycle makes and food processing machine countries of origin by region. Bajaj is the predominant motorcycle brand across the country, yet urban areas generally have a wider distribution of makes. For food processing machines, Kampala and Luwero firms mostly use Ugandan made machines, while Napak firms mostly employ Chinese made ones.

[32] Table A.4 in the appendix prints the results of Figure 1c in regression form and also reports estimates from the same specification with the paid repair price as dependent variable, which follow a similar yet noisy pattern. Unlike common machine models, the price paid to mechanics for repairing a rare brand is higher in remote areas, although the estimates are imprecise and not statistically significant.

[33] The finding that it can often take a long time to get rare or sophisticated machinery back operational after breakdown mirrors similar evidence found for government service delivery and infrastructure. Lee et al. (2020) for example report how a large electrification campaign in rural Kenya was hampered by persistent mechanical failures and large barriers to access adequate repair.

[34] Appendix C.5 provides details on the construction of market access measures for each parish in Uganda. There are only 29 mechanics in Napak district (21 for motorcycles and eight for food processing machines). Figure A.11 replicates all analyses of Figure 2 splitting along districts. Confidence intervals on the estimates from the most sparsely populated district are wide and cover both other districts.

Figure 2: Evidence on mechanic behavior by market access



(a) Time to serve arriving customers

(b) Motor mechanics skills test score

(c) Brands of spare parts in store

(d) Sourcing spare parts cited as problem

Notes: Evidence from the survey of mechanics investigating the mechanisms behind high replacement investment in remote areas. 95% confidence intervals around point estimates. Sample split by mechanic location market access quintiles, see appendix section C.5 for details on construction of this measure. In each panel, the set of leftmost estimates print the raw data without fixed effects. Panels a and d then introduce sector fixed effects (motor vehicle mechanics or food processing mechanics) and lastly firm size fixed effects (in quartiles). Panels b and c are only available for motor vehicle mechanics and so suppress the middle set of fixed effects. Time to serve arriving customers in panel a is computed as the stated time until a newly arriving order could be serviced, divided by the number of orders ahead in line. Skills test results in panel b show the fraction of answers in a non-incentivized skills test respondents got correct. Brands of spare parts in store in c show the number of motorcycle brand spare parts a mechanic currently has in store, out of the 18 most common ones (each mechanic was only asked a random subset of 3 brands). Panel d prints the share of mechanics listing "difficulty to find spare parts" as one of their three most pressing barriers to enterprise growth. Figure A.11 in the appendix replicates this analysis but splitting along study districts instead of market access quintiles. Because of limited power in Napak district (the survey team interviewed only 29 mechanics in this district), many gradients become undetectable in this specification.

to service a customer currently waiting for their machines to be repaired. This measure is constructed by eliciting the number of clients a mechanic currently has agreed to perform a repair task for, divided by a self-reported time estimate for when they could begin work on a new customer's repair order. In remote areas this processing time is around three times as long as in central ones, 114 minutes per task compared to 37. This pattern is robust to controlling for mechanic sector (motorcycles or food processing) and the number of employees a repair firm has.

Two further pieces of evidence offer insights into why mechanics are slower to respond to tasks in remote areas. First, central motorcycle mechanics perform around 20 percentage points better on a standardized skills test eliciting repair-specific knowledge (Figure 2b). This questionnaire was adapted with permission from Alfonsi et al. (2020), who developed it in collaboration with Ugandan vocational train-

ing institutes, and kindly agreed to share their survey materials. Respondents in central locations get about 65% of questions correct, compared to 48% in remote areas. Second, mechanics in well-connected parishes report to have around twice as many brand varieties of spare parts currently in stock (Figure 2c). Mechanics in the most central regions have spare parts for about six different motorcycle varieties in stock, compared to just over two in the most remote areas.[35] Lastly, difficulty in obtaining the necessary spare parts is also subjectively cited as one of the three most pressing issues inhibiting business expansion by around 60% of mechanics in remote areas, compared to only 20% in well-connected parishes (Figure 2d). This evidence is in line with remote areas having fewer motorcycle varieties on the road, rendering the owners of rare brands particularly vulnerable to cumbersome repair, as discussed above.

Together, the empirical patterns presented in this section point to an understanding of replacement investment as an endogenous object shaped by local economic conditions. Remote and small firms face additional difficulty to keep their capital intact. As these features are typically associated with firms in developing countries, these findings might generalize beyond Uganda, suggesting that the macroeconomic depreciation rate in poor countries might be larger than previously thought. I next present a model of the repair market which rationalizes these forces, and embedded in dynamic spatial equilibrium offers new insights into the interplay between development, depreciation, and growth.

4 A model of endogenous capital depreciation and economic growth

In this section, I present a model of capital repair and its contribution to economic growth. The key innovation of the model is that capital depreciation – usually treated as an exogeneous and constant parameter – is instead the endogenous outcome of a market process between machine owners and a sector supplying repair. I model this market and show that if it works poorly, the equilibrium replacement investment rate will be high, which lowers investment and ultimately throttles growth.

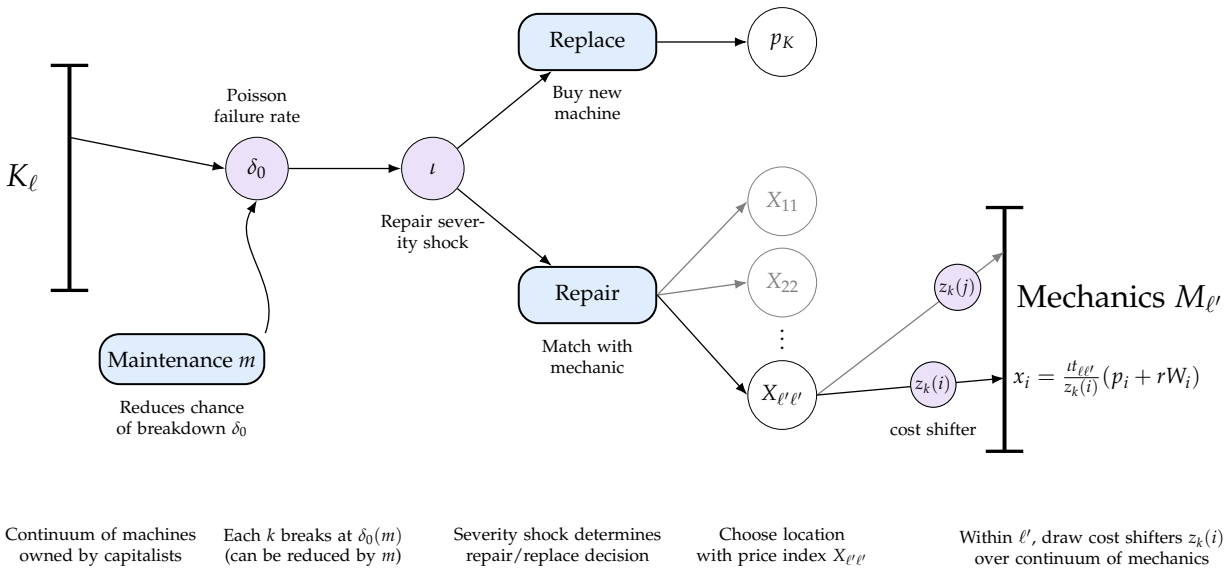
I first present a new microfoundation rationalizing why the costs to upkeep capital decrease with local economic development. I build on the insight that the timing of repair demand is unpredictable and timely, which I show gives rise to a non-standard benefit to agglomeration. Intuitively, if machines randomly break from time to time, mechanics in capital-scarce regions face periods of hardly anyone needing repairs followed by sudden spikes. Repair demand in capital-rich regions, however, approaches its long-run mean which mechanics can accommodate more easily, making it cheaper to upkeep capital if there is a lot of it. The model formalizes this intuition and derives tractable expressions for the resulting endogenous price and replacement investment rate. I first present these dynamics in a single-location benchmark, before embedding the framework in a general equilibrium spatial growth model. All derivations are in appendix section D, Figure 3 provides an overview of the model setup.

4.1 Model environment

Setup. Time is continuous and infinite. There are three sectors, subsistence agriculture, final goods production, and machine repair. Subsistence agriculture produces output Z using a decreasing returns to scale production function with labor N_Z but without capital. Firms produce final goods Y_G using

[35] In this survey module, respondent firms were asked whether they had three randomly drawn varieties (drawn from the 18 most common motorcycle brands in Uganda) currently in store. I construct the measure in panel 2c by multiplying the share of varieties respondents report having by 18.

Figure 3: Model setup



Notes: Visualization of the model setup. In each location ℓ, the representative capitalist owns a continuum of machines K_ℓ, which each fail at Poisson rate $\delta_0(m)$. Capitalists can perform maintenance m to bring down this rate. Once a machine has broken, capitalists draw a repair severity shock $\iota \sim \text{Exp}(\gamma)$ and decide whether to repair their machine or replace it with a new one (at price p_K). If they repair, they choose a location from which to source a mechanic, taking into account travel costs $t_{\ell\ell'}$ and local "factory gate" prices $X_{\ell'\ell'}$. Within each location, there is a mass $M_{\ell'}$ of mechanics indexed by i, who each provide mechanic services at a price p_i and with an expected wait time W_i. Capitalists additionally draw a cost shifter $z_k(i)$ for capital-mechanic pairs according to a Poisson point process defined in appendix section D.2 and choose the conditionally cheapest one.

capital K and labor N_G with a Cobb-Douglas production function:

$$Z = \Lambda N_Z^\beta, \qquad Y_G = AK^\alpha N_G^{1-\alpha}$$

Λ and A denote total factor productivity (TFP) in agriculture and final goods, respectively. Production in both sectors is perfectly competitive, final goods are sold at price p_G, and agricultural output is sold at an exogenous world price, which I normalize to 1. Labor markets and capital rental markets are also perfectly competitive, with labor paid a flow wage ω and capital a flow rental rate r. Workers are free to move between sectors. Consumers value consumption from both sectors according to

$$U = Z^\varphi G^{1-\varphi}$$

Capital accumulation. There is a continuum of machines forming the total capital stock K. I assume that capital is brittle and requires replacement investment at flow rate Δ, endogeneized below. Capital is owned by a unit mass of capitalists, who enjoy its rental income but are also responsible for its repairs, maintenance, and so on – in short: its replacement investment. New capital can be purchased at exogenous price p_K from abroad. Capitalists' investment and savings decisions maximize their intertemporal utility subject to a budget constraint:

$$v_s = \int_s^\infty \exp(-\rho t)\log c_t dt \qquad \text{s.t.} \qquad \underbrace{p_K \dot{K}_t}_{\text{new investment}} = \underbrace{r_t K_t}_{\text{capital income}} - \underbrace{\Delta_t K_t}_{\text{replacement investment}} - \underbrace{p_{C,t} c_t}_{\text{consumption}} \tag{2}$$

with ρ denoting capitalists' discount rate and p_C denoting the price of consumption (defined below). Workers cannot save and only consume their labor income hand-to-mouth (as in Kaplan and Violante,

2014; Kleinman et al., 2023). I show in appendix section D.1 that this problem is solved by a constant consumption rule $c_t^*(K_t)$ as in Moll (2014), leading to a steady state capital stock K^*:

$$p_{C,t}c_t^*(K) = \rho p_K K_t \qquad \text{and} \qquad K^* = \left(\frac{\alpha A p_G}{\Delta + p_K \rho}\right)^{\frac{1}{1-\alpha}} N_G \tag{3}$$

Replacement investment. As in section 3, the replacement investment rate Δ is the sum of all expenses a capitalist makes to keep a piece of capital in constant condition. Again, I assume that these expenses come in the form of direct repair costs, preventative maintenance, forgone revenue (in the form of lost rental income), and outright replacement. Capitalists optimally choose the relative proportions of each of these components to minimize overall replacement investment costs.[36] In particular, I assume each unit of capital breaks at an *iid* Poisson rate δ_0 at which point it becomes unusable until it gets repaired (at cost X, endogenized below) or replaced with a new unit from abroad (at cost p_K like all new machines). Unusable capital cannot be rented out, and so capitalists lose rental income while they wait to be served by a mechanic to repair their machine. For simplicity, I assume that purchasing a new unit is instantaneous and no forgone revenue accrues in this case. Lastly, capitalists optimally invest in preventative maintenance to lower the risk of breakdown δ_0 in the first place.

Combining these sub-problems, the key quantity determining the relative size of each of these components (and thus the overall replacement investment rate Δ) is the cost to repair a broken piece of capital, denoted by X. If X is high, all other components will be high as well – more is spent on maintenance to avoid full breakdown, and the chance that machines are replaced outright increases as repair gets more expensive. I first model this price X as the result of a matching process between machines and mechanics, before deriving the link between X and Δ below. I lastly characterize mechanic behavior, which is the key source of returns to scale in the model.

4.2 Repair demand

Repairing machines. To have their machines repaired, capitalists rely on a mass M of mechanic firms offering differentiated varieties of mechanic services. A given mechanic firm i is characterized by the price they charge for fixing a machine p_i, as well as the time this process will take at their store, which I assume is an exponentially distributed random variable with mean W_i. I, furthermore, assume that capitalists draw two sets of cost-shifters. First, a match-specific cost shifter $z_k(i)$ is drawn from a Poisson point process (discussed below) for every $k \times i$ machine-mechanic pairing once and stays constant throughout time. This shifter is meant to capture forces outside of the model that might influence which mechanic fixes which machine, such as having specialized knowledge in a certain machine type, or other more general trust and relationship benefits between particular mechanics and capitalists. Second, a breakdown-specific "severity" shock ι is drawn from an exponential distribution with parameter γ, $\iota \sim \text{Exp}(\gamma)$ every time a machine breaks. High realizations of ι represent difficult repair tasks which all mechanics find harder to fix. The severity shock is the same across mechanics and so does not influence which mechanic i fixes which machine k, but will influence whether the machine is repaired in the first place, with particularly tricky ι-draws being too costly to repair and instead resulting in replacement. Importantly, both $z_k(i)$ and ι are assumed to shift both the price p_i and the forgone rental income r in a symmetric way. To be precise, if mechanic i fixes machine k, the price accruing to capitalists is

[36] I abstract from input decay in the model, which represents only a small contribution to the overall replacement investment rate for machines in Uganda (see Table 2). The model could be adapted to include this in a straightforward way.

$\iota p_i / z_k(i)$, and they lose $\iota r / z_k(i)$ in rental income every instant the machine is broken and waiting to be repaired.

Capitalists hence value mechanics who are cheap and fast. I next make an important restriction on the timing of events. In particular, I assume that capitalists decide which machine k to entrust with which mechanic i by entering a *long-term* contract based on the expected average processing time W_i and the charged price p_i, not based on its realization at the moment of machine failure. In other words, capitalists make an *ex-ante* plan about which mechanic should fix which machine in case of breakdown, and do not deviate from it once that moment has arrived. This is important because the actually realized processing time at each mechanic is a random variable: there are times in which a mechanic is more or less busy and processing times adjust accordingly. I hence rule out a situation in which upon machine breakdown, capitalists renege on their ex-ante plan and instead condition the mechanic choice on whoever is free at that moment.[37] In turn, I assume that mechanics do not condition their price on the realization of their current level of demand, but instead keep it constant throughout.

Given this setup, capitalists decide where to send which machine in case of breakdown to maximize total net rental income less repair costs. This assignment problem is high-dimensional, as there are in principle infinitely many mechanics who could fix any given machine. To make progress, I follow Dupuy and Galichon (2014) and make the assumption that for each machine k, capitalists only consider a random but countable subset of "plausible" mechanics S_k. The remaining mechanics are assumed to be irrelevant for this particular machine, as they are unqualified to fix it. These mechanics might, however, be very well in the plausible set for a different unit of capital k'. For all mechanics within the plausible set for a given machine, I follow Dupuy and Galichon (2014) and assume that the cost shifters over mechanics $z_k(i)$ follow a Poisson point process with intensity $dx \times \frac{\sigma-1}{M} z^{-\sigma}$, where $x \equiv p + rW$ for readability, σ is a distributional parameter, and M is the mass of mechanics. As stated and discussed fully as assumption A.1 in appendix D.2, this implies that the probability that for a given machine k, capitalists have a mechanic with characteristics in the small neighborhood dx, dz around x, z is equal to $\frac{\sigma-1}{M} z^{-\sigma} dx dz$. Assumption 1 below summarizes this demand structure.

Assumption 1 *(Repair Demand) For each machine k, capitalists choose a mechanic i from the plausible set S_k based on the match-specific price $\iota p_i / z_i(k)$ and long-term expected wait time W_i, for which they forgo match-specific revenue $\iota r / z_k(i)$. Neither capitalists nor mechanics condition their choices on the current realization of the number of outstanding orders.*

As I show in appendix section D.2, this setup can microfound a solution to the machine×mechanic assignment problem which takes the shape of a continuous CES cost minimization problem for capitalists. In particular, I establish the following proposition:

Proposition 1 *Capitalists optimally allocating machines to mechanics on average incur repair and forgone rental income costs of*

$$\tilde{X} = \iota M^{\frac{1}{\sigma-1}} \left(\int_i (p_i + rW_i)^{1-\sigma} di \right)^{\frac{1}{1-\sigma}} \equiv \iota X \tag{4}$$

where σ is a parameter of the distribution of the cost-shifters $z_k(i)$ and M is the mass of mechanics. Similarly, the

[37] This assumption is in line with the high share of repeat interactions between mechanics and capital owners in the data: 90% of times a machine breaks, it is repaired by a mechanic the capital owner had previously done business with. Trust frictions plausibly contribute to this high number, as 65% of respondents agree that "most mechanics cannot be trusted".

Poisson rate λ_i at which new repair orders arrive at a given mechanic i is

$$\lambda_i = \zeta(p_i + rW_i)^{-\sigma} \tag{5}$$

where ζ is an endogenous market-level demand shifter ensuring repair market clearing, derived in appendix D.2.

Proof: see appendix D.2.

Proposition 1 restates the problem of where to send which machine in case of breakage as a tractable CES cost-minimization problem with elasticity σ. In contrast to conventional constant-elasticity solutions, the price index X aggregates over the sum of both the price p_i as well as the expected forgone rental income rW_i at each mechanic. Intuitively, mechanics will be allocated more market share if they are either cheap (low p_i) or fast (low W_i). Furthermore, as a result of the presented distribution of the cost shifters $z_k(i)$, the price index is also normalized by the mass of available mechanics $M^{\frac{1}{\sigma-1}}$. This is not necessary for the derivation to hold, but is done to turn off market-level love-of-variety effects and instead isolate returns to scale operating through the unpredictability of repair demand, introduced below.[38] Equation 5 translates this intuition into a residual mechanic demand curve. Importantly, note that this curve describes residual demand in *rates* of incoming repair orders. Because individual breakdown is a random event (happening at rate δ_0) and each mechanic is ex-ante assigned to a certain set of machines, so is the rate at which new orders come in. Cheaper or faster mechanics will garner more market share in the form of an increased – yet nevertheless random – Poisson rate of demand, as they cannot foresee exactly when their clients are next in need of a machine repair.

Replacing capital. Instead of repairing a machine every time it breaks, capital owners might also decide to scrap machines altogether and replace them with new pieces of equipment. Because I assume a newly repaired unit of capital is indistinguishable from a brand new one, the only reason to replace a machine is if it is cheaper than repairing it. This will happen for breakdowns with a severe shock ι, representing issues that all mechanics find difficult to repair. As discussed above, I parameterize the severity draws as an exponential distribution with parameter γ: $\iota \sim \text{Exp}(\gamma)$. This implies that the decision whether to repair or replace a machine can be summarized as:

$$Pr(\text{repair}) = Pr(\iota X \leq p_K) = 1 - \exp\left(-\gamma p_K / X\right)$$

Similarly, denote by $\mathcal{R}$ the expected cost of replenishing capital either through repair or replacement conditional on the above choice:[39]

$$\mathcal{R}(X) = \mathbb{E}_\iota \min[p_K; \iota] = \frac{X}{\gamma}\left(1 - \exp\left(-\gamma \frac{p_K}{X}\right)\right) \tag{6}$$

[38] A contribution of my paper is to present a new form of agglomeration effects which are distinct from more conventional love-of-variety effects. These effects could straightforwardly be added back to the model, and to the extent that they plausibly influence the market for repairs, the model as presented acts as a lower bound for the true amount of agglomeration one would expect for machine repair in Uganda.

[39] Two further useful object coming out of this parameterization are the rate at which machines are broken *and* subsequently repaired (instead of replaced):

$$\widehat{\delta_0} \equiv \delta_0 \left(1 - \exp\left(-\gamma p_K / X\right)\right)$$

and the unconditional expected value of ι for these cases:

$$\Psi \equiv \mathbb{E}_\iota[\iota] = \frac{1}{\gamma}\left(1 - \exp(-\gamma p_K / X)\left(1 + \gamma p_K / X\right)\right)$$

See appendix section D.3 for the full derivation.

and note that this is a monotonically increasing function of X: if repair costs are high, firms are forced to replace their machines for less severe breakages, which in turn increases expected costs. In the limit of extremely high repair costs, $\lim_{X\to\infty} \mathcal{R}(X) = p_K$, so intuitively machines get replaced every time they break down.

Preventative maintenance. Capitalists also invest in preventative maintenance to reduce the rate δ_0 of capital breaking down in the first place. I assume a simple functional form whereby maintenance m can be purchased at price c_m at the local mechanic, and the failure rate is parameterized as[40]

$$\delta_0(m) = \frac{1}{\nu} m^{-\nu}$$

Capitalists know from above that each time their machine does break generates costs of $\mathcal{R}(X)$ and optimally choose how much maintenance to purchase to minimize total costs. I show in appendix D.4 that this leads to a simple maintenance rule resulting in an optimal failure rate δ_0^* implying a total replacement investment rate (including breakdown and maintenance) of

$$\Delta = (1+\nu)\delta_0^* \mathcal{R}(X) \tag{7}$$

which crucially depends on how expensive it is for capitalists to get a machine repaired once it is broken.

The takeaway from these derivations is that the cost of repairs X is the critical object driving the size of all components of the replacement investment rate Δ. Next, I model mechanic behavior and derive X as the endogenous outcome of the market for repairs. I then discuss its implications for economic growth.

4.3 Repair supply

Repair is provided by a mechanic sector, which is subject to the key new force of the model: scale effects due to unpredictability of the timing of demand.

Repair production and unpredictability. The mass of mechanics M stands in monopolistic competition so that each mechanic firm seeks to maximize profits subject to their residual demand curve derived in (5). Once a machine has broken and a mechanic has begun working on it, repair is produced using a single variable input: spare parts s, which are available at exogeneous import-price p_S.

In order to be attractive to their customers, however, mechanics also need to offer fast service. I assume that mechanics can increase the expected speed at which they repair machines by hiring more labor n, available at market wage ω. In particular, n workers can solve incoming repair tasks at Poisson rate μn, where μ denotes the productivity of repair workers. Note that μn is itself a Poisson *rate*, as some repairs randomly take longer than others.

The key modeling choice of the market for capital repair is that demand is at the same time *unpredictable* and *timely*. Capital owners cannot pre-schedule *when* they need repair services, because machine breakdown is an inherently random event (it happens at rate δ_0). Once their machine does break, however, they value fast service without much disruption and lost rental income. Mechanics compete to increase the long run average rate at which demand materializes $\lambda_i(x_i)$, but cannot foresee the exact timing of when the next customer comes through the door. At the same time, I assume that hired repair labor

[40] For simplicity, I assume maintenance orders do not interfere with the timing of repair provided by local mechanics and are sold at marginal costs so do not generate profits for mechanic firms.

needs to be present continuously, it cannot be contracted on an order-by-order basis.[41] Because of this, mechanic labor might sit idle for periods of time with low demand, while times of excess demand lead to queuing and potentially long waits for customers. This presents a tradeoff for mechanics: choosing a faster rate μn by hiring more workers offers an attractively fast service for which they can charge higher prices. On the other hand, those workers might sit idle for stretches of no demand.

Demand of Poisson rate λ implies the time between two customers coming into a repair shop needing their machine fixed is exponentially distributed with mean $1/\lambda$, while supply of Poisson rate μn implies the time between finishing two successive repair jobs is exponentially distributed with mean $1/(\mu n)$. Under conditions of perfectly predictable demand and no costs to timeliness, the mechanic trivially would hire precisely enough labor to solve all incoming tasks, so that $n\mu = \lambda$. However, with the exact timing of demand being unpredictable, it might happen that an order comes in while a mechanic is still occupied with a previous one. In this case, I assume that a queue forms which a mechanic will attend to in a first-come-first-served manner. Recall that by assumption 1, customers do not balk at long queues but only adjust their assignment of which mechanic to entrust their machines with based on its long-run average. Assumption 2 below summarizes this supply structure.

Assumption 2 *(Repair Supply) A mechanic with n employees can solve repair tasks at Poisson rate $n\mu$ in order of arrival. Repair employees are not productive outside of repair production and receive their wage regardless.*

The setup described here is equivalent to a simple M/M/1 queue with demand rate λ and service rate $n\mu$, which is well understood in queuing theory.[42] A standard result is that the processing time for an incoming order between entering the queue, potentially waiting in line, and finally being processed is an exponential variable with mean

$$W = \frac{1}{n\mu - \lambda} \tag{8}$$

for $n\mu > \lambda$. I replicate a common proof of (8) in appendix section D.5. Note that this processing time goes to infinity as $n\mu$ approaches λ. This is because in the absence of spare capacity, slight perturbations in the exact timing of demand lead to a queue that piles up and can never be recovered. A mechanic with this waiting time is essentially useless for clients, who would incur large costs while waiting for their machine to be repaired. Mechanics hence hire more repair workers than strictly needed to have buffer capacity in case of sudden spikes in demand.

Mechanic problem. Mechanics decide how many workers to hire and how much to charge for their repair services. The central tradeoff is that hiring more workers than strictly necessary drastically reduces the expected time in which customers can be served, leading to more demand and the ability to charge higher prices. At the same time, these workers might sit idle in-between orders. The full

[41] This implies a mechanic cannot perform a different activity, say farming, and only switch to becoming a mechanic worker once a repair order has come in. Over 80% of mechanics in the Ugandan sample have not had any economic activity other than repairs in the last 6 months, despite the prevalence of low capacity utilization rates (see Table A.1 in the Appendix).

[42] In queuing theory, the M/M/1 shorthand denotes queuing processes that have a *M*arkovian or memoryless arrival time, a *M*arkovian service time, and a single server. Note that even though mechanics are increasing their service capacity by hiring more employees, I assume that all employees work together on a single task and do not form multiple service points. This assumption is made for model brevity and represents a lower bound on the processing speed benefits from adding employees to a queuing process.

mechanic problem can be stated as

$$\max_{p_i, n_i} \quad \lambda_i(x_i)\left(p_i - p_S\right) - \omega n \qquad \text{s.t.} \qquad x_i = p_i + \frac{r}{n\mu - \lambda_i(x_i)} \quad \text{and} \quad \lambda_i(x_i) = \zeta x_i^{-\sigma}$$

where the constraints follow from equations (5) and (8). I can now state a key proposition:

Proposition 2 *The optimal number of repair workers hired n^* is a concave function of repair demand λ:*

$$n^* = \frac{\lambda}{\mu} + \sqrt{\frac{\lambda r}{\omega \mu}} \tag{9}$$

As a result, expected wait time and mechanic marginal costs decrease with demand:

$$rW = \sqrt{\frac{r\omega}{\mu\lambda}} \qquad \textit{and} \qquad mc(\lambda) = p_S + \frac{\omega}{\mu} + 0.5\sqrt{\frac{r\omega}{\mu\lambda}} \tag{10}$$

Proof: See appendix D.6.

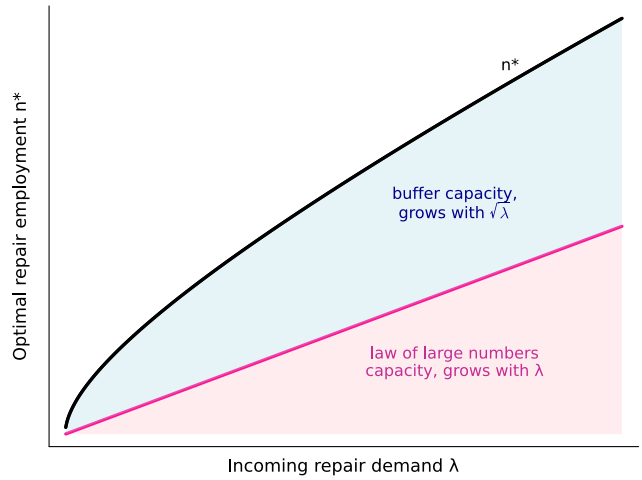
This relationship is the central source of scale effects in the repair market. Note how the first term of the optimal number of employees in (9) corresponds to the amount hired under perfectly predictable demand: each worker can service μ tasks, so λ/μ workers are needed for a total demand of λ. The second term $\sqrt{\frac{\lambda r}{\omega\mu}}$ corresponds to the buffer capacity that is hired to be immune to unpredictable spikes in demand. This buffer is larger if wages ω or productivity μ are low and if customers incur higher costs during waiting r. Crucially, it also grows with the square root of λ, which implies that it is relatively larger at low levels of demand. As λ grows, this term becomes less and less important as the law of large numbers begins to smooth out demand close to its long-run mean. In the limit, the relative importance of the term vanishes. As a result, the expected waiting costs incurred by a capitalist, as well as the marginal cost incurred by a mechanic to serve one additional expected client both *decrease* if a mechanic has more demand λ.

Figure 4 illustrates this intuition. Panel 4a plots the level of incoming repair demand λ and resulting optimal mechanic capacity choice n^*. This choice n^* is split into its two components, the capacity employed under perfectly predictable demand (or "law of large numbers capacity" in pink) and the buffer capacity to optimally smooth out demand shocks (in blue). For low levels of demand, mechanics need to employ relatively larger amounts of buffer capacity by hiring more workers than they strictly need.[43] The dark blue curve in Figure 4b prints the resulting marginal costs, which decline rapidly with λ and approach the limit of fully predictable demand $mc = p_S + \omega/\mu$.[44] Given this cost structure, monopolistically competitive firms charge an optimal markup over marginal costs, which also declines with overall demand λ. Panel 4b prints resulting prices in light blue, which approach the standard CES markup $\sigma/(\sigma-1)$ over marginal costs as λ grows, yet also are elevated considerably if demand is low. In general, no closed-form expression for the price p is available, yet it can be readily solved with available numerical methods. I, furthermore, show in appendix D.7 how to derive an analytical solution of the full model for the empirically relevant case of $\sigma = 4$. Considering the additional costs to capital owners from waiting, mechanics with low demand are even more unappealing as the composite

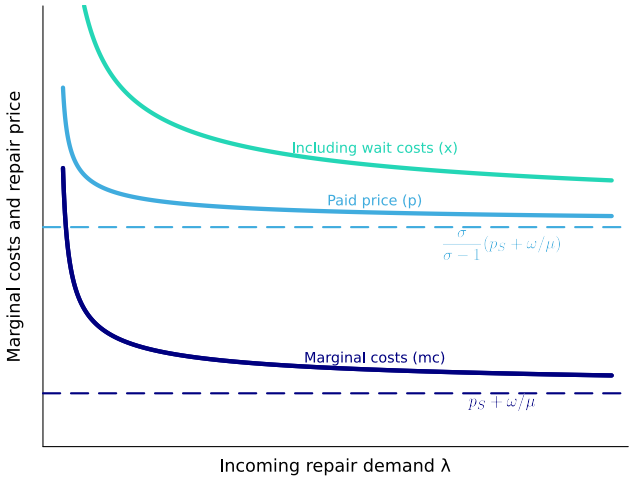
[43] In *absolute* terms, the buffer capacity of course still grows as demand increases (with $\sqrt{\lambda}$). The *relative* amount needed per unit of output, however, goes down rapidly. Figure A.14 in the appendix presents an alternative visualization of the dynamics underlying Figure 4a by printing the buffer and "law of large numbers" capacity as relative shares. For low levels of λ, the blue buffer component dominates.

[44] To see that $p_S + \omega/\mu$ are the marginal costs under fully predictable demand, note that each unit of repair needs one spare part at cost p_S and on average $1/\mu$ units of labor.

Figure 4: Unpredictability of demand leads to increasing returns to scale



(a) Optimal capacity choice n^* as function of λ



(b) Resulting costs and monopolistic prices

Notes: This Figure illustrates how the unpredictability of demand can lead to increasing returns in the supply of mechanic services. Panel a reports optimal firm behavior as a function of the long-run arrival rate of demand λ. The optimal labor choice of mechanics n^* (and thereby their resulting speed of service delivery) is a concave function of λ: $n^* = \lambda/\mu + \sqrt{\lambda r/\mu\omega}$. This is due to the additional capacity required to insure the mechanic against sudden demand spikes, which grows with the square root of λ (in blue). Panel b prints the resulting marginal costs to the mechanic of servicing one more client (in dark blue), as well as the resulting optimally charged price p_i and composite price x_i. Both of these decrease with demand, encapsulating the intuition that high demand allows mechanics to service everyone at cheaper costs. In the limit of infinite demand ($\lambda \to \infty$), marginal costs approach the costs under perfectly predictable demand ($p_S + \omega/\mu$) and prices the CES markup over them (dark and light blue dashed lines).

costs $x = p + rW$ approach infinity for low levels of λ (green curve).

Entry. To enter the repair market, I assume that potential mechanics incur a fixed flow entry cost f_E, denoted in units of labor. As in Melitz (2003), free entry implies that mechanics will enter until the profits from becoming a mechanic equal this fixed cost: $\pi_M = f_E\omega$.

Intuition and discussion. To build intuition for these results, consider a mechanic in a remote area in which there are only few machines that could break and hence customers arrive only rarely, say once a year. If the mechanic knew for certain that every year on the first of June the new order came in and there were no benefits to being fast, it could tailor its processing capacity perfectly to be able to perform exactly this one task every year. In practice, however, such a service would be worthless to customers, for two reasons: for one, this minimum viable repair capacity would take a full year to repair the machine, leaving businesses with idle for a very long time (demand is *timely*).[45] An even more pressing second issue appears if demand is not perfectly *predictable*. If new orders do not arrive in neat one-year intervals but rather at any point throughout the year, the chance that a new order arrives while the mechanic is still busy with the previous one is high. As orders pile up, the expected wait time will go to infinity, rendering the mechanic uncompetitive in the market. To avoid this, the mechanic is forced to hire costly buffer capacity which sits idle for much of the year, leading to high costs at low rates of demand. If the rate of incoming demand λ goes up so that on average every hour a new order

[45]This demand externality mirrors an effect first described by Mohring (1972) in the context of public transit provision. If hardly anybody rides the bus on a given line, the optimal amount of buses allocated to it implies service frequencies that are too rare to be attractive to riders, exacerbating the problem. On the other hand, very busy lines are allocated many buses, which hence arrive every few minutes, attracting even more riders, and so on. The same intuition holds in the context of repairs, another timely service: if many clients demand their machines repaired, mechanics devote additional capacity to solving tasks fast, rendering the resulting product more appealing to customers.

comes in, every order by itself does not present a big shock to the mechanic firm. Fluctuations around the long-run mean are small, and relatively less buffer capacity is needed. In the limit, the law of large numbers ensures that demand is perfectly predictable, allowing mechanics to employ vanishingly little relative buffer capacity and offer fast service at low costs.[46]

Unpredictability of demand hence gives rise to a non-standard form of increasing returns which operate at the mechanic firm level. Individual mechanics face lower average and marginal costs if demand is high, as fluctuations around its long run mean become less severe. These forces are distinct from a more conventional "market thickness externality" (Diamond, 1982; Miyauchi, 2024), in which the presence of *other firms* surrounding a business lowers its marginal costs. They are also different from cost curves derived under the presence of *fixed costs* or *input indivisibilities*, which would lower average – not marginal – costs with rising demand (Bassi et al., 2022). Continuously falling marginal costs also present a strong consolidation incentive to mechanics: it is cheaper for a single mechanic to serve 100 customers throughout the year than the combined cost for two firms serving 50 customers each, who have to deal with more idiosyncratic demand fluctuations. The CES structure of the demand system resulting from match-specific cost shifters $z_k(i)$ counteracts these incentives as mechanics have some monopoly power over their competitors with worse machine-specific draws (who might be less skilled in repairing this particular unit of capital). Monopolistically competitive mechanics restrict their own supplied quantity to maximize profits and new mechanic entrants will hence always attract some positive demand.

The question of whether scale effects could be achieved without any form of indivisibility in input choices was subject of a lively debate in the middle of the 20th century, with Lerner (1944) arguing perfect divisibility in inputs necessarily implies constant returns to scale and Chamberlin (1948) providing counterexamples. Whitin and Peston (1954) later argued that random fluctuations in demand, similar to the forces proposed here, can lead to increasing returns. Levhari and Sheshinski (1970) and later Arrow et al. (1972) connected this debate to the problem of optimally repairing machines *within* firms (known as the *repairperson problem* after Feller, 1957). Using a queuing theory setup, they prove that as firms get large, their optimal choice of capital and repair personnel approaches constant proportions, similar to the right tail of the results presented here. My model emphasizes instead that environments of low demand might be far away from satisfying these asymptotic conditions. In particular in capital-scarce and developing economy settings, machines just might not break frequently enough to escape the region of high marginal costs, low mechanic utilization, and high resulting replacement investment.[47]

As written, the model allows mechanic firms to increase the rate at which they repair machines by hiring more employees n. Made in the interest of model parsimony, this can be understood as a shorthand encapsulating various conceivable ways in which mechanics can increase their operating capacity. Other equally relevant margins of adjustment, such as having more spare parts in stock or hiring higher-skilled employees, would give rise to a similar tradeoff.

Taking stock, this structure of the repair market rationalizes high repair costs X in low-capital settings. This in turn implies that capital owners invest more in maintenance to avoid breakdown in the first

[46] An equivalent intuition holds for other interpretations of the optimal capacity choice, such as the stocking of enough spare part varieties. For areas or machine models with little repair demand, mechanics would have to keep spare parts stocked for a long and costly time before someone needs them. In regimes of frequent demand, the law of large numbers ensures that the probability that any spare part takes up shelve space for long becomes negligible, lowering costs and resulting prices (Arrow et al., 1951; Amberg et al., 2025).

[47] Recall from Table A.1 that the food processing mechanics have on average two customers a day, with a high week-to-week revenue ratio of almost 4.

place and are more likely to replace their machine with a new unit once it does break down as repair has become uneconomical. In short, replacement investment rates Δ are higher in areas with little repair demand. Such scale effects are hence a plausible driver of the stylized facts presented above.

4.4 Spatial economy

I next nest the discussed forces in spatial equilibrium to endogenize the level of repair demand and speak to the stylized fact that more developed regions have cheaper and faster repairs.

The economy is made up of L locations, indexed by ℓ. Each location is endowed with an immobile working population N_ℓ, two TFP terms for final goods and subsistence agriculture A_ℓ and Λ_ℓ, as well as a state variable capturing its capital stock K_ℓ. I assume new capital can be purchased at the same price p_K in all locations.[48] I also follow the assumption in Kleinman et al. (2023) that capitalists are immobile and cannot invest beyond their home location. Workers are assumed immobile. Space enters in two ways: firms can turn to mechanics in different locations and final goods can be traded across regions. I introduce each of these in order.

Final goods trade. Each location produces its own differentiated variety of the final good and consumers everywhere enjoy having access to more varieties. In particular, consumers in location ℓ consume the consumption bundle c_ℓ and derive indirect utility u_ℓ according to

$$c_\ell = \left(\sum_{j=1}^{L} c_{\ell j}^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}, \qquad \text{and} \qquad u_\ell = \frac{\omega_\ell}{p_{C,\ell}} = \frac{\omega_\ell}{\left(\sum_{j=1}^{N} p_{\ell j}^{1-\theta} \right)^{\frac{1}{1-\theta}}}$$

where $c_{\ell j}$ is the amount of location j's variety consumed by a consumer in location ℓ, $p_{C,\ell}$ is the consumer price index in location ℓ, $p_{\ell j}$ is the price of location j's variety in location ℓ, and θ is the elasticity of substitution. Final goods trade is subject to iceberg trade costs, so that $\tau_{j\ell} \geq 1$ units need to leave location j for one unit to arrive in ℓ, implying $p_{\ell j} = p_j \tau_{j\ell}$. Lastly, the share of ℓ's total consumption coming from location j is characterized by

$$\Theta_{\ell j} = \frac{(p_j \tau_{j\ell})^{-\theta}}{\sum_n (p_n \tau_{n\ell})^{-\theta}} \tag{11}$$

Mechanic trade. Machines in Uganda often get repaired by mechanics who have traveled far distances. To allow for this, I assume repair firms in each location produce a differentiated bundle of repair services at local price $X_{\ell\ell}$, where the double-subscript indicates the "factory gate" price of the repair bundle from location ℓ in its home location. Similar to final goods trade, mechanic travel is subject to iceberg trade costs, so that $t_{j\ell} \geq 1$ units of repair services need to leave location j for one unit to arrive in ℓ. Capitalists deciding who should repair their machine observe location-specific prices around them and maximize a second CES aggregator with elasticity ξ, leading to mechanic import shares $\Xi_{\ell j}$ and resulting price indices X_ℓ of

$$\Xi_{\ell j} = \frac{(X_{jj} t_{j\ell})^{-\xi}}{\sum_n (X_{nn} t_{n\ell})^{-\xi}}, \qquad \text{and} \qquad X_\ell = \left(\sum_n (X_{nn} t_{n\ell})^{1-\xi} \right)^{\frac{1}{1-\xi}} \tag{12}$$

[48] I do not find evidence of substantial trade barriers to purchasing *new* equipment in my data. If anything, observationally similar machines are cheaper in more remote areas.

where $\Xi_{\ell j}$ is the share of all machines from ℓ repaired by mechanics from j, and X_ℓ is the aggregate repair price index facing capitalists in location ℓ.

4.5 Equilibrium

In each location, total disposable income consists of labor income in all three sectors, agricultural profits, and capitalists' not-reinvested rental income:[49,50]

$$I_\ell \equiv \underbrace{\omega_\ell N_\ell}_{\text{labor income}} + \underbrace{\rho p_K K_\ell^*}_{\text{capitalists' income}} + \underbrace{\Pi_{Z,\ell}}_{\text{subsistence ag. income}}$$

Fraction $(1-\varphi)$ of I_ℓ is spent on final goods G from various locations, the rest on the freely traded and homogeneous subsistence agriculture good. In equilibrium, workers are indifferent between working in final goods, subsistence agriculture, or for a mechanic firm. Mechanics enter freely and make zero profits, and the labor market, goods market, and market for mechanic services clears in every location. I, furthermore, restrict attention to steady-state equilibria, in which capitalists save and consume so that the total stock of capital in each location remains unchanged. I define equilibrium as a collection of final goods prices p_ℓ, wages ω_ℓ, mechanic prices $X_{\ell\ell}$, rental rates r_ℓ mechanic masses M_ℓ, and final goods employment $N_{G,\ell}$ in all locations, such that mechanics and firms maximize profits and in all locations

1. Final goods markets clear: $p_\ell^c Y_\ell = (1-\varphi)\sum_j \Theta_{j\ell} I_j$
2. Repair markets clear:[51] $\int_{M_\ell} \lambda_i di = \widehat{\delta}_0 K_\ell$
3. Labor markets clear: $M_\ell N_\ell^* + N_{G,\ell} + N_{X,\ell} = \bar{N}_\ell$
4. Capital markets clear: $r_\ell K_\ell^* = \frac{\alpha}{1-\alpha}\omega_\ell N_{G,\ell}$
5. Free entry in the repair sector: $\pi_{M,\ell} = f_E \omega_\ell$
6. The economy is in steady state: $\rho p_K = r_\ell - \Delta_\ell$

I next use my detailed survey data on repair market dynamics across Uganda, combined with administrative microdata to calibrate my model to the Ugandan economic geography.

5 Estimation and calibration

I calibrate the model using a series of approaches. The model requires a set of novel structural parameters that have not been quantified in the literature to date, such as the travel costs faced by mechanics or the speed at which repair services can be provided. I harness my rich microdata on the demand and supply of capital repair across the country to identify and structurally estimate these parameters. I also draw on the existing literature to calibrate more established quantities. Table 3 presents an overview

[49] Using the decreasing returns to scale assumption in subsistence agriculture, one can express profits in this sector in any location ℓ as

$$\Pi_{Z,\ell} = \Lambda_\ell^{\frac{1}{1-\beta}} \omega_\ell^{\frac{\beta}{1-\beta}} \left(\beta^{\frac{\beta}{1-\beta}} - \beta^{\frac{1}{1-\beta}}\right)$$

[50] Also note that mechanics make zero expected profits after subtracting the fixed entry cost paid to labor.

[51] As defined in footnote 39 above, $\widehat{\delta}_0 \equiv \delta_0 \left(1 - \exp\left(-\gamma p_K / X\right)\right)$ denotes the arrival rate of machine failures that result in repair.

of the various estimates and their respective sources of identification. While the model is cast in continuous time, I express all temporal parameters as annualized rates. When taking averages across both surveyed sectors, I weight observations according to survey sampling weights.

Economic environment. To fit my model to Uganda's economic geography, I consider an adapted version of the country's administrative map. Uganda's first administrative subdivision are 138 districts, some of which are very sparsely populated and have limited data availability. To make progress, I consider a coarser geography in which some of the least populated districts are combined to form 51 larger units. For expositional clarity, I refer to the resulting locations uniformly as "districts" throughout, even though some locations are the amalgamation of multiple districts.[52] Wherever finer administrative data is available, I appropriately aggregate up to the newly coarsened level. Appendix Figure A.17 prints a map of the resulting geography.

For each district, I obtain the labor force and employment share in subsistence agriculture from the 2014 census. I also construct a measure of the total formal-sector value added in each location as well as the full formal-sector district-to-district trade matrix using VAT data for 2021 from the Ugandan Revenue Authority (URA). This dataset by definition includes the output of all formal firms in any district. While it misses all informal activity, it likely captures the majority of final goods trade and is estimated to cover around 50% of Uganda's GDP (Henning and Okello, 2024). I furthermore use data on the universe of customs records from the URA in 2022 to tally up total capital machinery imports into each district. Lastly, I use OpenStreetMaps to obtain a measure of travel distances and durations between all districts' largest towns.

Calibrated parameters $p_K, p_S, \alpha, \delta_0, \mu, \nu, \beta, \rho$**.** To calibrate the price of new machinery p_K, I use a combination of customs records and my survey data. In particular, capital goods imported into Uganda in 2022 had an average border price of \$976.[53] This value is by construction exclusive of tariffs and potential intermediary markups, which I approximate using my survey data on the average price paid by motorcycle riders for their machine, divided by the average import price of motorcycles at the border, or 1.45.[54] Using the assumption that the markup from border price to the price eventually faced by the capital owner is similar for motorcycles and all other capital classes, I arrive at $p_K = 1.45 \times 976 \approx 1417$. Spare parts for an average repair service in the survey cost around 2.5% of the price of a new machine. Again assuming this ratio is representative for other capital classes, I calibrate $p_S = 0.025 \times p_K = 35.3$. In my data, the average frequency at which machines are repaired is around 1.04 times per year. I add to this the reported flow rate of 0.26 that a machine is broken and not repaired to end up at $\delta_0 = 1.04 + 0.26 = 1.30$. To calibrate the elasticity of the maintenance function ν, I use the procedure outlined in appendix section D.4 and divide total replacement investment inclusive and exclusive of maintenance, to arrive at $\nu = 0.51$. Lastly, I assume that mechanic productivity μ_ℓ scales with final goods TFP A_ℓ between regions. To calibrate the overall level of repair productivity, I use my survey which asked mechanics the speed at which they typically repair a machine (not counting travel times). This averages to around 194 repair tasks per year per employee, which I use to arrive at $\mu_\ell = 194 \times A_\ell / \bar{A}$,

[52] As part of this coarsening, the study district Napak is joined together with four neighboring districts. The other two study regions remain their own distinct location in this simplified geography.

[53] This is the average per-item value of all capital imports into Uganda, belonging to HS code 84-90 (Caselli and Wilson, 2004). I exclude household electronics and a small number of HS4 codes with highly itemized parts, such as small batteries or diodes, which have a per-item value of just a few cents.

[54] Motorcycle imports into Uganda face a 15% tariff, leaving an additional 30 percentage points for wholeseller markups, well in line with existing estimates in the literature (Atkin and Donaldson, 2015; Grant and Startz, 2020).

Table 3: Model parameters

Parameter	Source	Value
New machine price p_K	firm survey & customs records	1,417
Spare part cost p_S	firm survey	35.3
Machine breakage rate δ_0	firm survey	1.30
Average mechanic repair rate μ	mechanic survey	194
Elasticity of maintenance function ν	firm survey	0.51
Capital share in final goods production α	ILO statistics	0.62
Agricultural consumption share φ	calibrate (Walker et al., 2024)	0.68
DRS parameter in subsistence ag. β	calibrate (Gollin and Udry, 2021)	0.43
Time preference parameter ρ	match Ugandan savings rate	0.11
Mechanic gravity parameter ξ	estimate on travel data	3.98
Mechanic travel costs parameter κ_m	gravity estimate on repair trade data	0.29
Final goods trade elasticity θ	calibrate (Atkin and Donaldson, 2022)	4.0
Final goods trade costs parameter κ_G	gravity estimate on VAT data	0.10
Mechanic elasticity of substitution σ	estimate jointly (markups)	4.16
Repair-replace parameter γ	estimate jointly (replacement rate)	0.29
Mechanic entry cost f_E	estimate jointly (capacity utilization)	3.34
Final goods TFP A_ℓ	model inversion (match trade shares)	51 values
Agricultural TFP Λ_ℓ	calibrate (Gollin and Udry, 2021)	51 values

Notes: Overview of model parameters and their source of identification. All temporal parameters correspond to annual quantities, monetary parameters expressed in US dollars.

where $\bar{A}$ is the average TFP across regions, derived below.

I use administrative statistics from various government agencies to calibrate α and ρ. According to the ILO, Uganda's labor share is about 37.7%, implying a capital share in final goods of about $\alpha = 0.623$. I calibrate the time preference parameter ρ using the Ramsey Rule and 2018 government statistics. Uganda's real lending rate (14.7%) and consumption growth rate (3.4%) together with the model-implied elasticity of marginal utility of one yield $\rho \approx 0.11$. Finally, to quantify the share of labor in the subsistence agriculture production function β, I rely on work by Gollin and Udry (2021), who using detailed agricultural panel data from Uganda estimate $\beta = 0.43$.

Partially relying on survey data to calibrate a model of capital repair covering the entire Ugandan economy requires making assumptions about the representativeness of the two surveyed sectors. To assess the degree to which the sectors are outliers in their relative repair burden, I compile administrative tax return data from the Ugandan Revenue Authority covering formal firms' expenditures throughout the Ugandan economy. The data allows me to construct a measure of how much firms spend on machine repair over the course of a year (an expense firms indicate on their annual corporate income statement), which I divide by the total book value of firms' machinery and equipment (as indicated on firms' balance sheets). Figure A.16 in the appendix prints the resulting distribution of repair expenditures across industries in the formal sector. Land transport and grain milling, the two ISIC-codes closest to the surveyed industries are no obvious outliers. While land transport does belong to the upper decile of industries in terms of how much formal firms report to spend on repairs, grain milling is in the lowest quintile, allaying concerns that the two surveyed industries are especially repair-intensive.[55]

[55]The measure discussed here differs from the first component of Δ introduced above in a number of ways. First, it does not include maintenance and forgone revenue expenditures, while replacement investment Δ does. These are particularly high for grain milling, partially explaining why the administrative repair shares are around half as large as the one I measured in the field. Second, Δ uses firm-appraised machine resale values as the numerator, while the administrative measure uses a book

The extent to which the more structural parameters of the model are representative presents a second challenge. This is particularly pressing for the annual rate at which machines break (δ_0), as well as the relative cost of spare parts (p_S), which might very well differ widely between machine types, capital classes, and locations (indeed, Table 1 shows that δ_0 varies between Ugandan districts by up to factor 2). In the absence of unified data on how often other machines break down per year, I make the admittedly unsatisfying assumption that my data is representative and $\delta_0 \approx 1.3$ is a reasonable estimate for capital throughout the Ugandan economy. Reassuringly, I do not find evidence of major variation in this rate across machine types *within my survey*: appendix Figure A.15a prints the breakdown rate for the nine most common machinery types (four different motorcycle brands and five different food processing machine types), as well as the engineers' estimate of farm machinery in the US (ASABE Standards, 2011), which are all in reasonable proximity to the calibrated estimate of $\delta_0 = 1.3$. A similar pattern is observed for the price of spare parts relative to new machinery (p_S/p_K, Figure A.15b). Future research collecting comprehensive evidence on how these parameters vary across further capital classes and machinery types would be very insightful.

Gravity parameters $\xi, \kappa_M, \theta, \kappa_G$. I calibrate the spatial and gravity parameters using my detailed survey data. I follow Ahlfeldt et al. (2015) and express iceberg trade costs for both mechanic (t_{ij}) and final goods trade (τ_{ij}) as

$$t_{ij} = \exp(\kappa_M \text{dur}_{ij}) \qquad \tau_{ij} = \exp(\kappa_G \text{dur}_{ij})$$

so that κ_G, κ_M capture the travel cost elasticity with respect to travel duration dur_{ij} for final goods and mechanics, respectively. I measure dur_{ij} as the travel duration between locations i and j as reported by OpenStreetMaps, in hours.

For mechanics, I rely on a survey module which asked repair firms how much they would charge if, hypothetically, they received an order from a customer at various, randomly drawn parts of the country. Taking logs, I can isolate κ_M using a simple regression

$$\log t_{ij} = \kappa_M \text{dur}_{ij} + f_i + f_j \tag{13}$$

where f_i and f_j are mechanic location and travel destination-level fixed effects. Table A.5 and Figure A.18 in the appendix print estimation results and scatter plots for all gravity regressions. Using the above approach, I obtain an estimate of $\kappa_M = 0.29$ (column 1). Next, following the approach by Ahlfeldt et al. (2015) and taking logs of equation (12), I can use my data on the origin districts i of mechanics servicing firms in a given district j, as well as the above estimate for κ_M, to identify the elasticity of substitution between mechanic origins ξ (defined in equation 12):

$$\log \Xi_{ij} = -\xi \kappa_M \text{dur}_{ij} + f_j \tag{14}$$

which yields $\xi = 3.98$, very close to the consensus estimate of trade elasticities around four (Table A.5, column 2). Note that equation (14) does not include mechanic-origin fixed effects f_i, which I suppress to gain power as many mechanic origin districts only feature once or twice in my sparse data.[56]

value depreciated according to the depreciation schedule. Third, the sample of firms is different, as formal land transport and grain milling firms are likely larger and concentrated in major cities.

[56] Column (3) of Table A.5 prints resulting estimates with the full set of fixed-effects. Standard errors in this underpowered regression are an order of magnitude larger than in the remaining specifications, leading to an implausibly large estimate of $\xi = 9.8$. I hence prefer the specification without the full set of origin fixed effects.

I use a similar strategy to uncover the parameter κ_G related to final goods trade. Without direct data on the cost components of final goods traders, I resort to the consensus estimate of $\theta = 4.0$ commonly used in the trade literature in African countries (as reviewed by Atkin and Donaldson, 2022). I then use VAT records from the Ugandan Revenue Authority capturing the district to district network of intranational formal trade to fit an equation similar to (14) by taking logs of equation (11):

$$\log \Theta_{ij} = -\theta \kappa_G \text{dist}_{ij} + f_i + f_j \tag{15}$$

Using this procedure, I obtain $\kappa_G \theta = 0.38$ or $\kappa_G = 0.10$ (column 4 of Table A.5).[57]

Structural parameters σ, γ, f_E. I jointly estimate the within-location elasticity of substitution between mechanics σ, the shape parameter of the repair-issue distribution γ, as well as the fixed costs of mechanic entry f_E using Simulated Minimum Distance. In practice, I assume the economy is in steady state, and for a candidate vector $\Omega = [\sigma, \gamma, f_E]$, I solve for the general spatial equilibrium outlined above, construct a series of identifying moments, and identify the values of Ω that exactly fit the true data moments. I use the set of moments

$$\begin{bmatrix} \sigma \\ \gamma \\ f_E \end{bmatrix} \rightarrow \begin{bmatrix} \text{Markup for mechanics over spare parts} \\ \text{Share of breakages resulting in replacement} \\ \text{Mechanic capacity utilization} \end{bmatrix}$$

Below, I outline a brief discussion of why these three moments plausibly identify the resulting set of parameters. Figure A.19 in the appendix additionally presents pseudo-"Jacobians", derived from making local changes in each parameter and re-computing the resulting moment vector.

Intuitively, the elasticity of substitution σ is a component of market power for monopolistically competitive mechanics, who can charge higher markups over the cost of a spare part if σ is low. I focus on the markup charged over spare part prices to avoid the issue of having to apportion wage payments (which are difficult to elicit in owner-operated enterprises) onto each individual repair task. In the data, the average price charged for a repair is about 59.5% above the cost of the used spare parts. In joint estimation, this is matched by a parameter of $\sigma = 4.16$, well within the consensus range of these estimates in the literature (Atkin and Donaldson, 2022).

The parameter γ governs the distribution of the severity draws ι capitalits draw each time a machine breaks. This intuitively impacts the tradeoff of whether to replace or repair, with a high value of γ implying more replacement as the chance of a particularly tricky severity draw becoming larger. In the data, firms repair a machine 79.9% of the times it breaks (and replace the rest of times), which is matched by an estimate of $\gamma = 0.29$.

Lastly, the labor-denoted entry fixed cost for mechanics f_E shifts the local density of mechanics relative to the amount of machinery in the market. If these fixed costs are high, few mechanics enter and so

[57] All regression estimates in this section are obtained using OLS. The quantitative spatial literature frequently estimates gravity equations like the above using PPML. However, as my data on mechanic travel flows are incomplete along one dimension (my survey captures all origins of mechanics traveling to the three survey locations but misses all other destinations of these mechanics to districts outside of the study area), PPML will spuriously interpret missing entries in the travel-flow matrix as evidence for high distaste of mechanic trade. Columns 5-6 of Table A.5 repeat the main gravity estimations using this procedure. Not surprisingly, the resulting estimate of ξ (column 5) is around twice as large as in the OLS specification (column 2), which is my preferred specification. The PPML estimate for final goods trade (which doesn't have this problem as VAT flows data is available for the full district network) is close to the OLS estimate. Both estimates for κ_M and κ_G are close to the value estimated by Walker et al. (2024) using a similar specification for local trade flows in Western Kenya ($\kappa = 0.22$)

each remaining one is left with relatively more demand. Building on this intuition, I identify f_E with the aggregate capacity utilization of mechanics, which is 58% in my survey. In joint estimation, a parameter of $f_E = 3.34$ matches this value, implying that in order to start a repair enterprise, potential entrants in the mechanic sector pay a fixed cost roughly three times the yearly market wage, or \$3,900.

Economic fundamentals A_ℓ, Λ_ℓ. To close the calibration exercise, I need to estimate total factor productivity in all locations for both subsistence agriculture Λ_ℓ and final goods A_ℓ. For subsistence agriculture, I rely on data on agricultural productivity collected by the World Bank Integrated Surveys on Agriculture. This nationally representative panel survey collected detailed measures of agricultural production Z_ℓ and labor inputs N_Z from about 5,000 households across Uganda between 2005 and 2015. This is the same data Gollin and Udry (2021) leverage to estimate the parameter β I use above. Since I assume land is fixed, I can directly back out implied TFP from output and labor choices for each household h: $\Lambda_{\ell,h} = Z_{\ell,h}/N_{Z,h}^{\beta}$. Taking the average across households in each district, I obtain an estimate for Λ_ℓ. Figure A.20a in the appendix prints the spatial distribution of both the underlying data as well as the implied productivity terms. Reassuringly, the coffee-growing regions in the center and south-west of the country display high values of Λ_ℓ.

Lastly, I back out final goods TFP parameters A_ℓ using model inversion. The conventional approach in the literature inverts the wage equation and given data on the distribution of labor and capital sets TFP to the vector rationalizing observed data on wages (Bilal and Rossi-Hansberg, 2023). This procedure critically relies on estimates of the capital stock K, which is usually derived using the perpetual inventory method using a fixed rate of depreciation δ, in conflict with the empirical findings presented above. I hence rely on a nested routine leveraging the final goods gravity equation (15). In particular, the fixed-effects f_j in (15) can be written as[58]

$$f_j = -\theta \log \Gamma_0 \omega_j^{1-\alpha} r_j^{\alpha} + \theta \log A_j \tag{16}$$

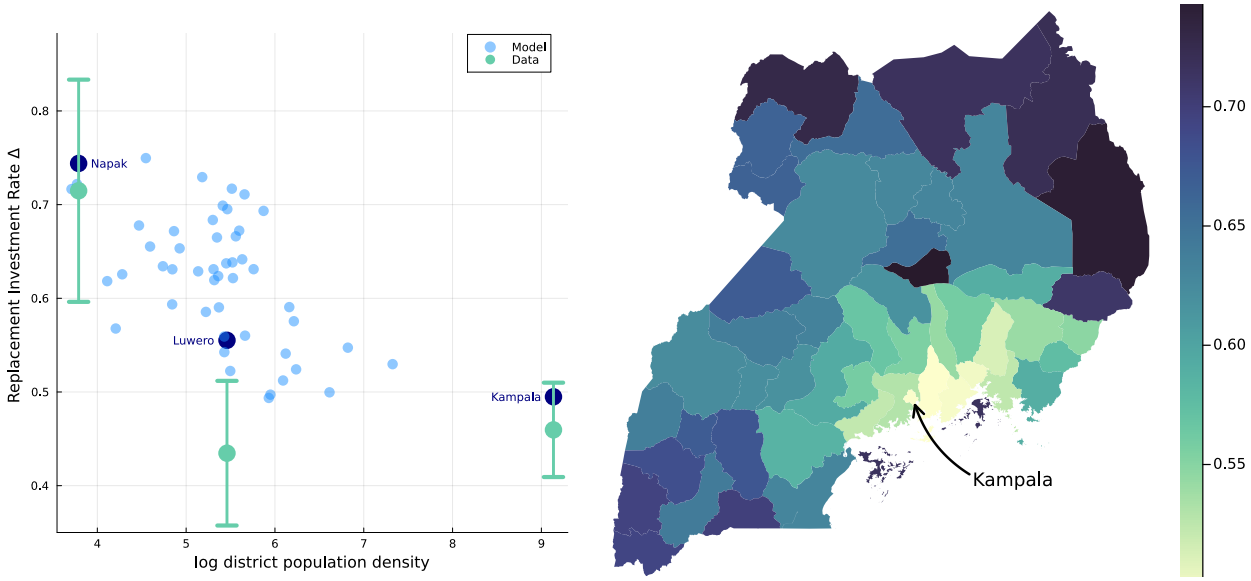
This suggests a simple algorithm, whereby for a given vector of structural parameters (σ, γ, f_E) and a guess of the A-vector, I compute the equilibrium, use the estimated f_j to update A_j according to (16), and iterate until convergence. Section E.1 explains this procedure in more detail. Figure A.20b prints the spatial distribution of TFP A, which is highest in Kampala and its surroundings.

Model fit. The fully estimated model matches a series of untargeted characteristics of the Ugandan economy. Figure A.21 in the appendix compares model predictions against administrative data on the distribution of capital and labor across Ugandan districts. Panel A.21a assesses model fit on the total amount of machinery imported into each region, data for which come from administrative customs records. I focus on annual capital *imports* and not estimates of the capital stock, as the latter is usually constructed using a constant and uniform accounting estimate of the depreciation rate. The model does a good job predicting the staggeringly wide distribution of local capital import intensities, which range from about 1000\$ of imported machinery per capita per year to just under 10cts.

Panels A.21b and A.21c compare model-predicted employment shares in repair occupations and subsistence agriculture against their data equivalent from the 2014 Census. The model matches the relative sizes of each local repair market decently well (although it predicts a wider distribution than the truth), and roughly identifies the most agriculturally-heavy regions (although not perfectly). Finally, panel

[58] $\Gamma_0 \equiv \left(\frac{\alpha}{1-\alpha}\right)^{1-\alpha} + \left(\frac{1-\alpha}{\alpha}\right)^{\alpha}$ is a constant.

Figure 5: Model-predicted replacement investment rate Δ across Uganda



(a) Δ against district population density

(b) Δ across Ugandan districts

Notes: Replacement investment rates Δ across regions in Uganda. Panel a on the left prints model-predicted Δ rates against the log population density (in blue). The three study regions are highlighted. Green estimates print the empirical Δ estimates from the firm survey with 95% confidence intervals. This is the same data as Figure 1a. Panel b on the right prints the same as a map, with the most remote regions in the north-east displaying the highest replacement investment rates.

A.21d compares GDP per capita in model and data. As sub-national GDP estimates are not published, I construct a measure of value added in the formal sector from value-added tax records in each district, which I then uniformly scale up to account for informal activity. While model and data generally agree in the largest regions, the tax records are much sparser in the poorest Ugandan districts, where informality plausibly is much higher, significantly skewing the spatial distribution of economic activity according to the tax data (Henning and Okello, 2024).

6 Endogenous depreciation and economic development

I next use the estimated model to investigate the macroeconomic implications of endogenous capital depreciation under scale effects. I first quantify the extent to which repair markets contribute to differences in capital intensity and economic activity across Uganda. I then simulate the effects of a series of counterfactual development policies aimed at increasing capital investment, before exploring two quantitative extensions to the baseline model.

6.1 Replacement investment and development accounting across Uganda

Scale effects, trade, and the geography of economic activity all influence how expensive it is to upkeep capital in the calibrated model environment. Figure 5 presents the predicted replacement investment rate Δ for all model regions plotted against district population density. Panel 5b prints the same model prediction on a map. The model successfully replicates the stylized facts #1 and #2 from section 3.3: Δ is high throughout Uganda and highest in remote regions. Indeed, the predicted replacement investment rate is about 50 percent in Kampala and 25 percentage points higher still in the most remote regions of the country. The Figure also prints the corresponding data estimates from the three study regions in green (these are the estimates from Figure 1a). The admittedly wide confidence intervals

cover the model prediction for Napak and Kampala districts, though the model slightly over-estimates the replacement investment rate for peripheral Luwero district.

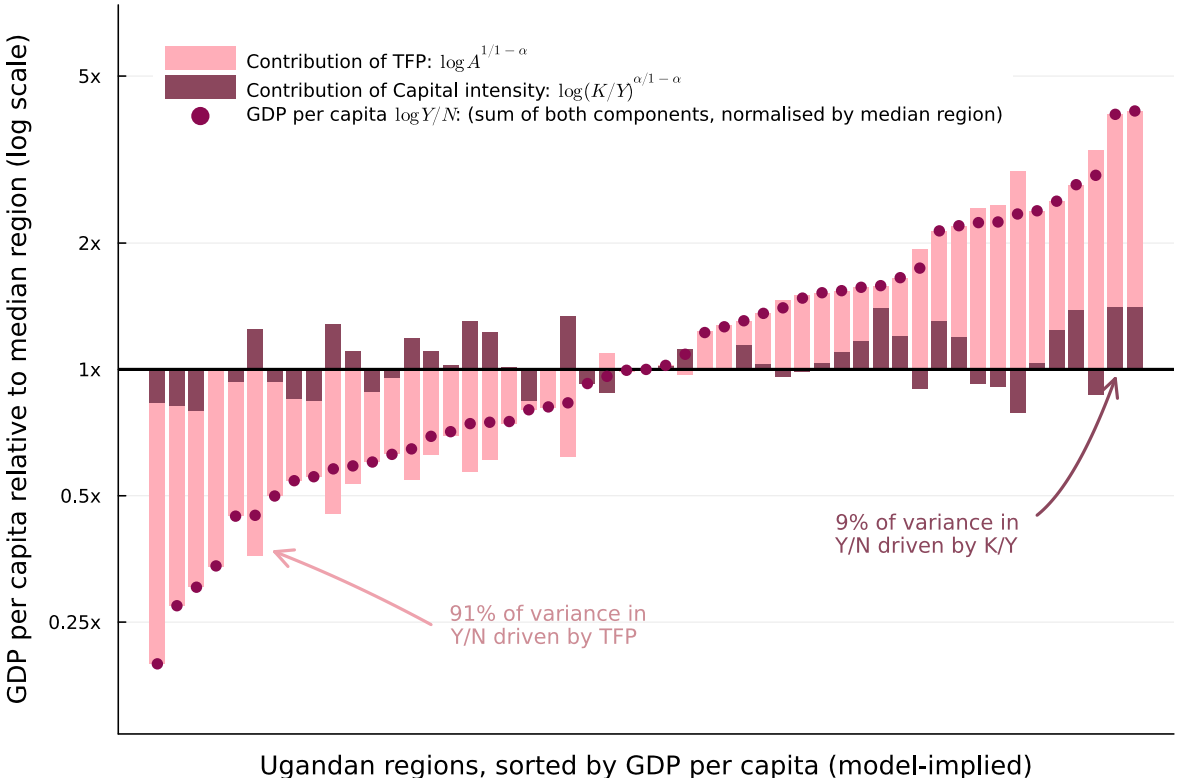
What are the implications of different depreciation rates for economic development? Intuitively, scale effects in the market for repairs make investing in the least developed regions more costly, leading to fewer machines being used in production, which further amplifies initial productivity differences. To quantify the size of this multiplier effect and its ramifications for income differences across space, I conduct a development accounting routine across regions in Uganda. The goal of this exercise is to assess how much of the differences in income per capita between places can be attributed to observable differences in factor inputs, and how much remains as a TFP or Solow-residual term. In particular, I follow Klenow and Rodriguez-Clare (1997); Hall and Jones (1999) and rewrite the final goods production function:

$$\begin{aligned} \frac{Y_\ell}{N_{Y,\ell}} &= A_\ell^{\frac{1}{1-\alpha}} \left(\frac{K_\ell}{Y_\ell} \right)^{\frac{\alpha}{1-\alpha}} \\ &= A_\ell^{\frac{1}{1-\alpha}} \left(\frac{\alpha}{\Delta_\ell + \rho p_K} \right)^{\frac{\alpha}{1-\alpha}} \end{aligned} \tag{17}$$

The first line is an accounting identity which decomposes final goods sector output per worker of each region ℓ into the relative contributions of capital intensity K_ℓ/Y_ℓ and TFP A_ℓ. The second line writes out the steady-state capital intensity and follows from the law of motion of capital (derivation in appendix section D.1). Two parameters in equation (17) vary between regions ℓ: the replacement investment rate (a high Δ_ℓ implies a low capital intensity and hence less output per capita) and the residual A_ℓ. Apportioning variations on the left hand side to these two components hence speaks to the relative importance of endogenous depreciation in explaining the differences in living standards across space. Below, I first conduct this exercise within the calibrated model environment across Ugandan regions, where I am directly able to derive the model-predicted Δ. In the absence of uniform and well-measured data on Δ for every region, this analysis is less a pure accounting exercise, and should rather be understood as assessing how much income variation the calibrated model can generate and attribute to the channel of endogenous depreciation. Section 7 presents more suggestive evidence extrapolating this mechanism beyond Uganda to explore how differences in depreciation might impact global differences in output per worker.

Figure 6 visualizes the accounting decomposition of equation (17) across Ugandan regions. Focus first on the dark dots, which represent model-implied output per worker in the final goods sector $(Y_{G,\ell}/N_{G,\ell})$ for all 51 regions in Uganda, sorted from poorest to richest. I exclude subsistence agriculture to focus on the sector subject to capital-accumulation and growth. To emphasize the variation in this measure, I normalize all local measures by the median region (without loss of generality) and plot differences on a log scale, so that the poorest (richest) region has less than a quarter (almost five times) the output per capita of the median one. Two remarks about this spread are in order. First, this more than 20$\times$ difference in output per worker between the richest and poorest region *within* one country is of a similar order of magnitude as world income differences *across* countries, a staggering but well-documented fact in global development (see e.g. Acemoglu and Dell, 2010). Second, in the absence of detailed subnational government statistics, Figure 6 uses model-implied estimates of productivity. As discussed above, the model matches the labor and capital distribution for which estimates do exist (see Figure A.21 in the appendix). The best estimate I can construct using administrative value-added tax data of non-agricultural formal sector GDP per capita covaries strongly with the model estimates but displays even more variation, likely due to differences in business formality (Figure A.21d). In light

Figure 6: Contribution of depreciation to income differences across Ugandan regions



Notes: Decomposition of model-implied output per worker in the final goods sector across Ugandan regions. Dark dots represent $\log Y/N$, normalized by the median region, and regions are sorted from poorest to richest (implying the richest region is almost $5\times$ as rich per capita as the median region, which is themselves more than $5\times$ richer than the poorest region). Light bars print the contribution of $\log A^{1/1-\alpha}$, dark bars the contribution of $\log(K/Y)^{\alpha/1-\alpha}$, each again normalized by the median region. The p-value on the slope of $\frac{\alpha}{1-\alpha}\log(K/Y)$ in this regression is $p < 0.01$. By accounting identity (17), the two components add up to overall output per worker. Variance decomposition applied via equation (18).

of this, the model-consistent estimates of output per worker are my preferred measure of development differences across regions in Uganda.

How much of this spread can be attributed to differences in depreciation? The dark and light red bars in Figure 6 decompose $\log Y/N$ into capital intensity $\log(K/Y)$ and the residual TFP $\log A$, appropriately rescaled according to equation (17). As everything is on a log-scale, both components add up to the dark dots and the relative size of each set of bars visualizes the relative contributions to overall income differences. Two things stand out. First, the light bars (corresponding to the residual TFP term $\frac{1}{1-\alpha}\log A$) are generally the dominant component for each Ugandan region, which means the majority of why some Ugandan regions are more productive than others remain unexplained even after accounting for differences in capital intensity due to differences in depreciation rates. Second, however, notice that the dark bars (corresponding to the capital intensity term $\frac{\alpha}{1-\alpha}\log(K/Y)$) covary with the object of interest: poorer regions tend to have a lower capital intensity than the median region (negative dark bars towards the left), while richer regions tend to be more capital intensive (positive dark bars towards the right). In particular, the richest region has about 42% higher capital intensity than the poorest region. A simple variance decomposition written as

$$\frac{cov(\log Y/N, \frac{\alpha}{1-\alpha}\log K/Y)}{var(\log Y/N)} \approx 0.089 \tag{18}$$

implies that variations in capital intensity can account for around 9% of the differences in output per

worker across Ugandan regions, while the other 91% remain unexplained as TFP. At the same time, there are some deviations from this trend, as some poorer regions have lower depreciation rates and hence higher capital intensities than the median region (and hence Figure 6 displays the occasional positive dark bar to the left of the median), and vice versa. On aggregate, however, the presented channel helps explain part of the large underlying spread in output per worker across space.

The magnitude of this decomposition warrants some discussion. Attributing 9% of differences in output per worker to differences in depreciation rates compares to active research on the role of schooling and human capital, which is commonly estimated at 10-30% (Hsieh and Klenow, 2010). On the one hand, my findings imply that if all regions in Uganda had the same depreciation rates, the difference in income per capita between the richest and the poorest would narrow from 20× to 18×.[59] On the other hand, by far the largest source of income differences remains TFP or "the measure of our ignorance" (Abramowitz, 1956). A further caveat to these findings is that a series of model ingredients make the environment *ex-ante* conducive to attributing capital with a big role in generating prosperity. For one, the Ugandan capital share as reported by the ILO of $\alpha = 0.623$ is large compared to rich countries, but not out of the ordinary for developing economies (Maarek and Orgiazzi, 2020). One concern with estimates of the capital share in poor countries is that they partially explain rents accruing to land, natural resources, and other non-reproducible capital (Caselli and Feyrer, 2007). In Uganda, 7.5% of GDP comes from natural resource rents, 5× more than the global average. As a robustness exercise, I adjust for this by calibrating $\alpha = 0.623 - 0.075 = 0.548$ and recomputing the equilibrium capital intensity and contribution to income differences. Making these adjustments reduces the share of total income variation attributed to endogenous depreciation to 7%.

A second concern with the above exercise is the use of model-implied data on the left-hand side. As discussed above, a measure I constructed from administrative tax records of output per worker displays a substantially larger spread between the richest and poorest Ugandan regions, which can be plausibly attributed in large parts to lower rates of formality and tax-compliance in the poorest areas. As a robustness exercise, I also interpret the tax data as a ground truth, which widens the variation to be explained and hence reduces the contribution of differences in depreciation rates to 3%. Given the issues with this subnational dataset and the fact the model fits well a series of more representatively estimated patterns of economic activity across Ugandan regions, my preferred estimates rely on the calibrated model, implying a contribution of endogenous capital depreciation to differences in prosperity of 7–9%.

6.2 Policy tools to increase capital investment

Given the importance of the repair market to regional growth and development, what tools does the government have to reduce depreciation rates and boost capital investment? In this section, I simulate the effects of various counterfactual policies and discuss which policy levers the model views as most effective to address an investment landscape characterized by endogenous depreciation. Figure 7 summarizes the main proposed policies.

Mechanic training. Repairing machines requires a host of sector-specific skills. Mechanics in Uganda acquire these skills through three main ways: 22% of survey respondents completed a certified vocational training program, 51% learned on-the-job as an apprentice in an existing repair firm, and the rest

[59] The differences in capital intensity between high and low depreciation areas are similar in magnitude to studies of the effects of capital market reforms in India (Bau and Matray, 2023).

are self-taught.[60] Given the high prevalence of firm-provided apprenticeships or on-the-job training, designing these programs in a way that incentivizes high-quality knowledge transmission has been the focus of an active research and policy debate in Uganda (Alfonsi et al., 2020). In a randomized controlled trial, Alfonsi et al. (2020) study the labor market effects of subsidizing firm-provided vocational training across eight sectors, which includes motor vehicle mechanics. In one part of their study, they estimate that the subsidy increased firm-provided training, which in turn increased apprentice productivity by 39% for workers in this sector over the control group.[61] I find similar differences in my observational data: mechanic respondents who have completed an apprenticeship score higher on the skills test and serve clients faster (though the latter is imprecisely estimated, see Figure A.12).

What would be the downstream effects on capital investment of scaling this program and increasing mechanic productivity throughout the country? To investigate this, I scale up mechanic productivity μ everywhere by 1.39 and recompute the steady state equilibrium. The leftmost panel of Figure 7 reports results. Making all mechanics in the country 39% more productive lowers the aggregate replacement investment rate Δ by two percentage points, from 0.51 to 0.49. In response, Uganda's aggregate capital intensity is predicted to increase by 3% and total output by 2%. Despite the increase in final goods production, the simulation predicts a slight uptick in subsistence employment, as making repair workers more productive renders some workers currently employed in this sector obsolete. In general, these effects are small given the large counterfactual increase in productivity. This is best understood as an application of the Baumol (1967) effect, whereby capital repair acts as a bottleneck technology in the aggregate economy. If labor in this sector is unproductive, costs are high and many resources are bound up. Once it becomes more and more productive, however, labor merely shifts away to the next bottleneck sector without major effects on aggregate GDP. Figure A.22 in the appendix illustrates this asymmetry by rerunning the above counterfactual for a range of mechanic productivity improvements. Relative to baseline, even a tripling of productivity would only decrease Δ by around four percentage points – while making mechanics less productive would have much more severe negative effects.[62]

Tariffs on spare parts. Uganda at present charges a 15% tariff on most imported machinery parts. I investigate the effects of removing this tariff. I assume full pass-through of the tariff onto mechanics' spare parts sourcing price p_S, but allow their markup charged to repair customers to adjust endogenously. I furthermore assume that only 62% of spare parts are imported, as the remaining machines in the sample are Ugandan-made. To ensure budget-balancedness, I impose a flat tax on total production to recoup the lost tariff revenue to the government. The second panel of Figure 7 shows that such a tariff reduction would be passed down onto the repair costs faced by firms, with Δ decreasing by about 3 percentage points. Uganda's capital intensity increases by more than 5%, total production by around 3% even after deducting the new flat tax.

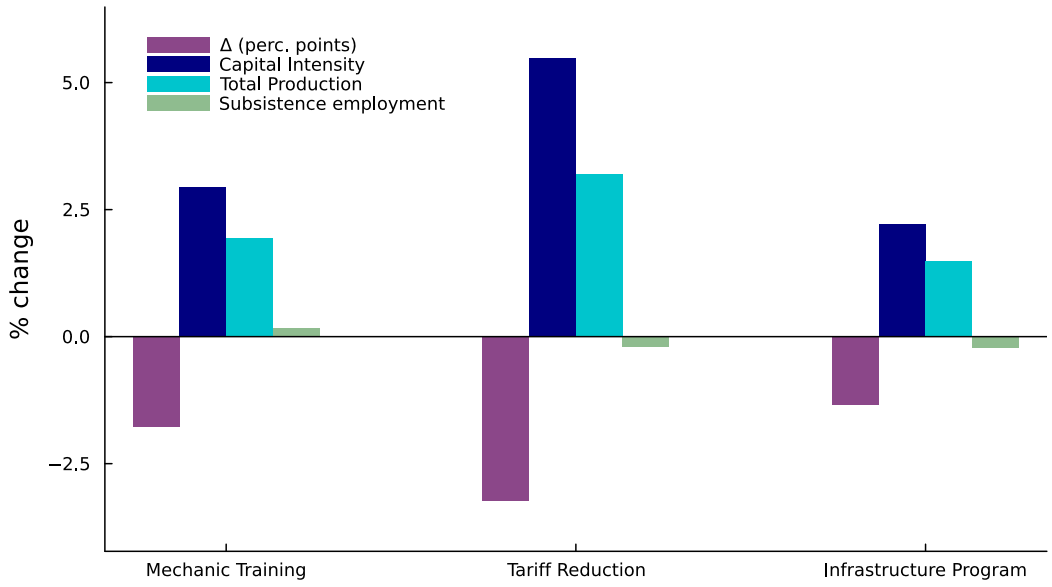
Infrastructure interventions. I lastly consider the effects of a hypothetical major infrastructure investment campaign which reduces travel times on all links by 10%. The cost to upkeep machinery responds

[60] 4% of mechanic respondents completed both an apprenticeship and a certified program. Food processing mechanics have higher rates of formal training: 33% compared to 12% for motorcycle mechanics. Apprenticeship rates are similar across sectors.

[61] In their paper, the authors only report the aggregate skill increases across sectors, which is around 32% compared to the control group (Alfonsi et al., 2020, Table 2, column 4). I rerun their replication file to isolate the effect for just the mechanic sector, which is slightly higher at 39%.

[62] Alfonsi et al. (2020) additionally conclude that despite these productivity increases to apprentices, costs to firm-owners providing this training are large enough to lead to an estimated negative internal rate of return for this subsidy program. The modest downstream effects on investment I estimate are likely not going to overturn this conclusion.

Figure 7: Counterfactual implications of policies intended to increase capital investment



Notes: Counterfactual predictions of movements in labor, capital, and repair markets in response to simulated policies. As explained in the main text, "Mechanic Training" increases the rate at which mechanics solve tasks μ, "Tariff Reduction" reduces the cost of spare parts p_S, "Infrastructure Program" reduces travel times dur_{ij} between all links by 10%. Δ changes are printed as percentage points, the remaining variables in percentages relative to baseline. Capital intensity and total production only refer to the final goods sector which uses capital, not to subsistence agriculture.

to this program through two channels: a direct effect making it cheaper to access mechanics from further away, and an indirect effect through increased overall production and repair market scale effects. In total, both effects together reduce Δ by about 1.3 percentage points and increases total production by 1.5%.[63]

6.3 Model extensions and robustness

6.3.1 In-house repair

As the repair market is plagued by high costs, firms might find it attractive to circumvent it by hiring their own mechanic staff. Indeed, stylized fact #3 above showed that firms with their own repair workers have lower replacement investment rates. So why don't all firms do this? One answer is that the very scale effects that made repair provision difficult in remote markets also make in-house repair costly for small firms. A microenterprise with just one or two machines does not generate enough demand to justify having someone on staff (and on the payroll) year-round waiting for the occasional breakdown. As firms grow and take on more capital, however, they generate a steadier flow of repair needs, making it more and more economical to take on in-house capacity. The model as written does not feature firms of different sizes and hence cannot directly speak to this dynamic.

In appendix section D.8, I sketch an extension of the baseline model with heterogeneous firms of varying productivity levels. Firms can decide whether they want to rent capital at a rate that includes replacement investment (as in the main model), or instead want to be responsible for supplying their own repair. In this case, they hire in-house mechanic workers subject to the same technological constraints summarized by assumptions 1 and 2: the timing of repair demand is unpredictable and workers need to

[63] While a flat 10% decrease in travel times is a coarse simulation of what an optimal infrastructure program could look like, the magnitude of these effects is within range of existing estimates in the literature. In Graff (2024), I estimate consumption increases from an optimal 10% expansion of Uganda's network to be around 2%.

be on the payroll even while they wait for the next breakdown. This leads to the same tradeoff as at the market level: hiring enough standby personnel to quickly respond to sudden repair needs is expensive if it is hardly ever needed. I formalize this intuition and show how the resulting equilibrium behavior is consistent with empirical stylized fact #3. Figure A.29a presents total capital costs for firms of varying productivity levels under both regimes, panel b prints resulting replacement investment rates. Small firms with few machines would have to hire an uneconomical amount of repair capacity and so rely on the expensive repair market, while larger firms optimally circumvent it leading to continuously lower replacement investment rates.

The insight that repair scale effects operate at both the market and firm-level also offers one suggestive explanation for the empirical regularity that market-based repair is so much more prevalent in poor countries (see Figure A.2b in the appendix). Where firms are small, there frequently is not enough capital to rationalize hiring a full-time employee tasked with repairing machines. The repair market steps in to fill this gap, yet is plagued by high markups. Together, these two effects exacerbate each other, rendering this problem particularly severe for the millions of capital-intensive microenterprises around the world.

6.3.2 Fixed repair productivity

The calibrated model successfully replicates stylized fact #2 that more capital-scarce regions have higher replacement investment rates. The main focus of the model have been costs of remoteness and scale effects in the repair market. However, my calibration also varies repair personnel productivity μ_ℓ with local TFP A_ℓ, leading to mechanics in Kampala to be about $14\times$ more productive than those in Napak (yet also earn a commensurately higher general equilibrium wage). To ensure that these parameterizations don't drive the proposed effect, I also re-calibrate the entire model as a robustness exercise while imposing μ to be fixed at its average of $\mu = 194$ across Uganda.

Figure A.23 in the appendix prints the resulting replacement investment rates under this specification against district population density, similar to Figure 5a. Predicted Δ values are lower throughout and the remoteness gradient is slightly less pronounced, yet the overall takeaway remains the same: the most remote districts have replacement investment rates around 25 percentage points below the most central ones, driven by scale effects in the repair market. Under this specification, the accounting contribution of model-predicted capital stocks rises slightly compared to the main model, to 14%. The main calibration hence generates varying Δ rates mainly because of trade costs and scale effects, not local mechanic productivity differences.

7 Global development accounting

Places where capital upkeep is more expensive likely have less of it. This applies to regions across Uganda, but also plausibly to global differences in capital intensities across countries. Indeed, existing international estimates of capital-output *ratios* (ie. the capital stock divided by annual GDP) differ by a factor of around four between the world's richest and poorest countries (Hsieh and Klenow, 2010). One challenge with this important fact in global development, however, is that the way capital is generally measured is sensitive to the choice of a depreciation parameter. How then do measures of the global capital distribution change once taking into account that capital depreciation might be higher in poorer regions? In this section, I discuss this problem before conducting a back-of-the-envelope reappraisal

of the level of K across the process of development and reassess a canonical development accounting exercise on the importance of capital as a source of global income differences.

The challenge with most commonly used international capital series is that they do not directly measure the stock K_t of machinery, vehicles, and buildings in a country each year. Instead, they combine data on capital investment I_t in a given year with an estimate of the depreciation rate and a base year K_0 to carry forward the previous year's capital stock according to the perpetual inventory method (PIM):

$$K_{t+1} = (1 - \delta)K_t + I_t \tag{19}$$

The choice of depreciation rates δ hence affects the measurement of capital stocks around the world. In practice, early development accounting exercises typically assumed a constant value of δ across space and time, often set to the US value.[64,65] More recent estimates have taken into account that the *composition* of capital varies between countries: poor countries have relatively more buildings (which depreciate slowly) while rich countries have relatively more high-tech equipment (which depreciates fast). The newest generations of the Penn World Table (Feenstra et al., 2015) account for these differences, which typically imply economy-wide average depreciation rates *rise* with development.[66] For any given capital class, however, adjustments for potential changes in the rate at which otherwise similar machinery is replaced across the development process are not generally performed.

The empirical evidence presented in section 3 above suggests that, at least within Uganda, motorcycles and food processing machines are replaced more quickly in less developed regions. The theoretical model in section 4 offered a microfoundation for how repair market imperfections can rationalize these patterns. Under the strong assumption that these trends generalize to other capital classes and beyond Uganda, I discuss below how accounting for these differences changes existing estimates of capital stocks around the world.

7.1 A back-of-the-envelope adjustment of K with varying depreciation

In this exercise, I recompute K for every country in the Penn World Table using the perpetual inventory method after adjusting δ based on interpolating income differences between Uganda and the US. I first present this approach, then discuss its undoubtedly strong assumptions, before reporting the resulting data series and the implications of this adjustment for global development accounting.

Approach. To construct an adjusted global capital stock series, I largely follow the most recent generation of the Penn World Table (PWT, Feenstra et al., 2015). PWT computes K_{it} for country i in year t using equation (19) and data on annual investment for up to 183 countries since 1950. Investment series I_{ait} are further split into four different asset classes a: buildings, machinery, vehicles, and other capital.[67] Each asset class comes with its own depreciation rate δ_a, which is assumed constant across space and time and identical to the US value used by the Bureau of Economic Analysis going back to Fraumeni

[64] Studies using a constant value for δ include Mankiw et al. (1992), Hall and Jones (1999), Acemoglu and Zilibotti (2001), Hsieh and Klenow (2007), and Caselli and Feyrer (2007), as well as early iterations of the Penn World Table.

[65] Furthermore, the base year K_0 is frequently derived using an estimate of δ, the growth rate g, and the assumption that the economy was in steady state in that year: $K_0 = I_0/(\delta + g)$ (see for example Hall and Jones, 1999).

[66] The International Monetary Fund applies a similar adjustment at least since 2010 (Arslanalp et al., 2010).

[67] In practice, PWT uses a more detailed set of nine capital subcategories (for example splitting buildings into non-residential and residential structures). Detailed investment series are only available for the above four coarse classes which are aggregates of the subgroups.

(1997). Table A.6 in the appendix prints these rates and compares them to my survey evidence. To adjust these δ_a values for the level of economic development, I make two strong assumptions:

Cross-Country Assumption 1 *The relationship between GDP per capita Y/N and depreciation rates δ_a for Uganda and the US linearly extrapolates to all other countries.*

Cross-Country Assumption 2 *Depreciation rates δ_a are correctly measured in the United States.*

Assumption 1 fixes the *slope* of how δ varies with economic development. Another way to interpret this assumption is that differences local GDP are the only factor driving different depreciation rates across the world, and that this mapping is identical everywhere. Assumption 2 fixes the *intercept* of the world-wide level of δ_a in each year. Fitting a line through the two points US and Uganda (which according to Table A.6 has roughly $1.5\times$ as high physical depreciation rates as the US), I arrive at an elasticity of δ_a of ≈ -0.12 to GDP per capita.[68] Using the second assumption, I then adjust each country's depreciation rates based on the difference of their GDP per capita to the United States, recompute an updated capital-stock according to equation (19), and iterate until present day.[69]

Discussion of assumptions. The above assumptions might fail in numerous ways. First, high replacement investment rates Δ, which have been the focus of this paper, do not automatically translate into high *physical* depreciation rates δ for the purposes of the perpetual inventory method. Intuitively, while Δ quantifies the costs to keeping capital constant, δ quantifies how much of the physical capital stock vanishes every period. This distinction is important as many components of Δ such as regular repairs and maintenance are not usually captured as investment in the National Accounts. It hence could be that developing countries have high replacement investment without discarding their capital at higher rates – there would then be no mismeasurement in global capital stocks. However, recall from Table 2 that Ugandan rates of physical machine replacement (the "replacement" component of Δ) for the machines I study are also strongly elevated compared to the US, at factor of almost $3\times$. The slope identified under assumption 1 compares this physical replacement component to the conventional rates used in the PWT, which comes out at $1.5\times$ and is hence at the lower end of elasticities one could choose.

Second, assumption 1 will fail outright due to both omitted variable bias (things other than Y/N influence depreciation rates across the world) or external validity concerns (the US-Uganda difference is not informative for other countries). These are important concerns, but in the absence of better data on global depreciation rates, no immediate remedy is available. The results below should hence be treated with caution and understood as merely a proof of concept of how large the resulting measurement error in K might be.

Results. With these important caveats in mind, Figure 8a presents the K-series adjustments for 178 countries. The x-axis prints log GDP per capita from the Penn World Table, the y-axis the adjustments made using the outlined procedure. By construction and assumption 2, the US does not see their capital stock updated through this procedure. In contrast, many poorer countries are predicted to have less

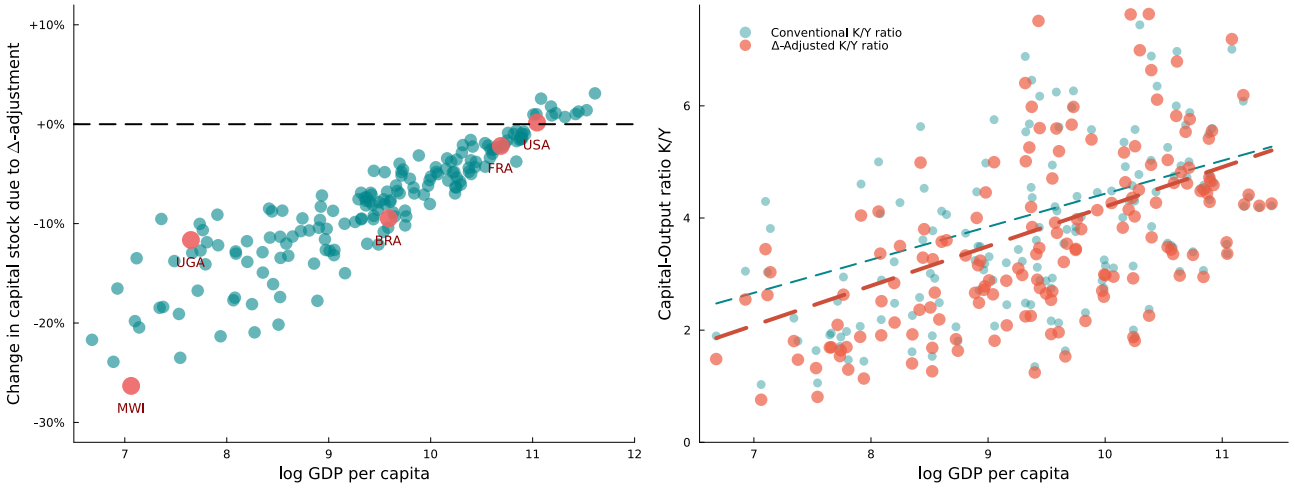
[68] The US had a roughly $31\times$ higher GDP per capita as Uganda in 2019, which yields $\log(1/1.5)/\log(31) \approx -0.12$.

[69] The exact adjustment performed for asset class a and country c is

$$\delta_{a,c} = \delta_{a,US} \times \left(\frac{Y_c/N_c}{Y_{US}/N_{US}} \right)^{-0.12}$$

which is performed for every year the PWT has data availability, and assuming that the capital stock in each country's base year is correctly measured.

Figure 8: Adjusting global capital stock measures for endogenous depreciation



(a) Global capital ratios after Δ-adjustment

(b) Adapted K/Y covaries more with prosperity

Notes: Panel a visualizes revisions made to global capital stocks from the Penn World Table (Feenstra et al., 2015) after accounting for depreciation of a given asset class being higher in poorer areas. It prints the relative ratio of the adjusted 2019 capital stock compared to the conventional measure on the y-axis, against 2019 log GDP per capita. The US serves as a benchmark and is not revised by assumption 2. Poorer countries are revised downwards by interpolating the empirical differences between Uganda and the US for every year since 1950. Panel b prints resulting capital-output ratios against log GDP per capita, both for the conventional capital series (in green) and the newly revised one (in orange), which displays a stronger comovement with GDP per capita (p-value of orange fit being steeper than green fit $p < 0.001$). Underlying these computations are strong assumptions which are discussed in the main text.

capital than implied by the traditional method. Intuitively, the high costs to repairing machinery and equipment means existing units of capital are more likely to fall into permanent disrepair than traditionally assumed. This in turn implies poor countries might have less capital than we think. The magnitude of this potential mismeasurement is substantial: Uganda is predicted to have around 10% less capital, Malawi over 25% less. On average, the poorest quarter of countries have around 15% less capital than traditionally measured. A few countries richer than the US are predicted to have the opposite mismeasurement, as scale effects are predicted to push depreciation rates below their traditionally assumed US values. This finding is heavily dependent on assumption 2 and hence even more speculative. Future research estimating depreciation and replacement investment rates in other advanced economies might discipline the magnitude of this direction of the mismeasurement of global capital stocks.[70]

7.2 Reassessing the importance of capital in comparative development

If poor countries have less capital than we think, the contribution of capital to global income differences might be higher than conventionally measured. Figure 8b visualizes this intuition. It prints in green how traditionally measured capital-output ratios covary with income per capita. As is well known, richer countries have more capital relative to output. Table 4 performs a development accounting exercise as in section 6 using equation (17). Using the conventional measures, column (1) attributes 20% of global income differences to differences in capital intensity, in line with existing studies such as Hsieh and Klenow (2010) (Figure A.25 in the appendix provides detailed, country-by-country contributions similar to Figure 6). The revised capital series printed in orange displays a stronger covariance with world income levels, as poorer countries have less capital than previously thought. This effect elevates

[70] Figure A.24 prints a histogram of all revisions in the depreciation rate under this technique. To minimize the impact of one-off outlier years, I cap the factor at which δ can be revised by $\pm 5\times$, which is rarely exceeded.

Table 4: Contribution of capital-intensity to global income differences

	(1) **PWT K-series**	(2) **Revised K-series**
Contribution of $\frac{\alpha}{1-\alpha}\log\left(\frac{K}{Y}\right)$	20%	25%
Contribution of $\frac{1}{1-\alpha}\log\left(A\right)$	80%	75%

Notes: Development accounting using equation (17) to apportion global income differences to capital-output ratios and residual TFP. Numbers indicate a covariance decomposition according to equation (18). The revised capital series in column (2) explains more of the global income variation because of higher depreciation rates in poor countries leading them to be revised to have less capital. Underlying these computations are strong assumptions which are discussed in the main text.

the estimated contribution of capital to global income differences to 25%, a roughly one quarter increase.

Summing up, the calibrated model of how repair markets lead to high replacement rates in capital-scarce regions hints at measurement error in the way capital stock series are usually constructed. A revised time series is about a quarter more predictive of global income differences than the conventional estimates. Because of the strong assumptions needed to arrive at this result, the implied magnitudes should undoubtedly be treated with caution. Adding first-hand evidence on replacement investment rates for firms across the world, similar to what this paper has attempted for Uganda, would allow for a richer understanding of the link between depreciation, growth, and development.

8 Conclusion

The rate at which capital breaks and is repaired is not an engineering constant, but varies across space and the process of development. Despite its importance for economic growth, this dynamic has received only limited empirical or theoretical attention. In this paper, I provide detailed evidence on the various ways in which the brittleness of capital impacts the operations of microenterprises across Uganda. I make four contributions. First, I introduce a new and direct way to group together these expenses into a comprehensive measure of replacement investment. Second, my firm survey isolates new stylised facts: capital replacement costs are high throughout Uganda, highest in remote locations and for small firms, and subject to scale effects. Third, I argue that a unique feature of the supply of repairs – the fact that they are impossible to schedule in advance – can explain these facts. I embed this feature into a general equilibrium spatial growth framework with endogenous depreciation and estimate it using my microdata. Fourth, I use the model to show that accounting for these effects can explain up to 9% of the variation of income across regions in Uganda, discuss the impacts of counterfactual investment policy tools, as well as offer a suggestive calculation of what these results imply for the measurement of capital across the world.

Understanding capital depreciation as an endogenous object subject to scale effects has implications beyond what is studied in this paper. For one, it offers a strong coordination incentive for firms to use similar machinery as their neighbors, offering a new microfoundation for why seemingly similar firms so often cluster close together in developing country settings (Bassi et al., 2022; Vitali, 2022; Weiss et al., 2024). Relatedly, my findings have potential implications for the dynamics of technology adoption in poor countries, as firms have an incentive not to be the first-mover to avoid paying excessively high repair costs. A vintage model in the spirit of Solow (1959) could operationalize this intuition and offer a quantification of the dynamic productivity implications of this effect (Caunedo and Keller, 2021).

An intriguing avenue for future work is to consider ways for governments to help firms overcome the challenge of accessing adequate and timely repairs. The technological challenge discussed in this paper – the fact that repair demand cannot be scheduled – is hard to tackle directly. However, because this problem becomes less severe in capital-dense areas and for large firms, well-agreed upon objectives of development policies might help alleviate these problems as well. Better integrated markets, a wider variety of skilled labor, and larger firm sizes all offer theoretical ways out of the low capital – high depreciation equilibrium.

References

Abramowitz, Moses (**1956**). Resource and Output Trends in the United States since 1870. *American Economic Review* 46(2), pp. 5–23

Acemoglu, Daron and Melissa Dell (**2010**). Productivity Differences Between and Within Countries. *American Economic Journal: Macroeconomics* 2(1), pp. 169–188

Acemoglu, Daron and Fabrizio Zilibotti (**2001**). Productivity Differences. *The Quarterly Journal of Economics* 116(2), pp. 563–606

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf (**2015**). The Economics of Density: Evidence From the Berlin Wall. *Econometrica* 83(6), pp. 2127–2189

Alfonsi, Livia, Oriana Bandiera, Vittorio Bassi, Robin Burgess, Imran Rasul, Munshi Sulaiman, and Anna Vitali (**2020**). Tackling Youth Unemployment: Evidence From a Labor Market Experiment in Uganda. *Econometrica* 88(6), pp. 2369–2414

Amberg, Niklas, Richard Friberg, and Chad Syverson (**2025**). Have We Got News For You: Firm-Level Evidence on the Optimal Choice of Expected Capacity Utilization. NBER Working Paper 33400, Cambridge, Mass

Arrow, Kenneth (**1963**). Uncertainty and the Welfare Economics of Medical Care. *American Economic Review* 53(5)

Arrow, Kenneth J, Theodore Harris, and Jacob Marschak (**1951**). Optimal Inventory Policy. *Econometrica* 19(3), p. 250. Place: Evanston

Arrow, Kenneth J., David Levhari, and Eytan Sheshinski (**1972**). A Production Function for the Repairman Problem. *The Review of Economic Studies* 39(3), p. 241

Arslanalp, Serkan, Fabian Bornhorst, Sanjeev Gupta, and Elsa Sze (**2010**). Public Capital and Growth. IMF Working Paper

ASABE Standards (**2011**). Agricultural Machinery Management Data

Atkin, David and Dave Donaldson (**2015**). Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs. NBER Working Paper 21439

——— (**2022**). The Role of Trade in Economic Development. *Handbook of International Economics* 5

Balboni, Clare, Oriana Bandiera, Robin Burgess, Maitreesh Ghatak, and Anton Heil (**2022**). Why Do People Stay Poor? *The Quarterly Journal of Economics* 137(2), pp. 785–844

Banerjee, Abhijit and Esther Duflo (**2005**). Growth Theory through the Lens of Development Economics. In: *Handbook of Economic Growth*, volume 1, pp. 473–552. Elsevier

Bari, Faisal, Kashif Malik, Muhammad Meki, and Simon Quinn (**2024**). Asset-based microfinance for microenterprises: Evidence from Pakistan. *American Economic Review* 114(2)

Barro, Robert J (**2021**). Double Counting of Investment. *The Economic Journal* 131(638), pp. 2333–2356

Barro, Robert J. and Xavier Sala-i Martin (**1992**). Convergence. *Journal of Political Economy* 100(2), pp. 223–251

Bassi, Vittorio, Jung Hyuk Lee, Alessandra Peter, Tommaso Porzio, Ritwika Sen, and Esau Tugume (**2023**). Self-employment within the firm

Bassi, Vittorio, Raffaela Muoio, Tommaso Porzio, Ritwika Sen, and Esau Tugume (**2022**). Achieving Scale Collectively. *Econometrica* 90(6), p. 53

Bau, Natalie and Adrien Matray (**2023**). Misallocation and Capital Market Integration: Evidence From India. *Econometrica* 91(1), pp. 67–106

Baumol, William J. (**1967**). Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis. *American Economic Review* 57(3), pp. 415–426

Bilal, Adrien and Esteban Rossi-Hansberg (**2023**). Anticipating Climate Change Across the United States. Working Paper

Bitros, George C. (**1976**). A Statistical Theory of Expenditures in Capital Maintenance and Repair. *Journal of Political Economy* 84(5), pp. 917–936. Place: CHICAGO

Blitz, Rudolph C. (**1958**). Capital Longevity and Economic Development. *American Economic Review* 48(3), pp. 313–329

——— (**1959**). Maintenance Costs and Economic Development. *Journal of Political Economy* 67(6), pp. 560–570

Brancaccio, Giulia, Myrto Kalouptsidi, and Theodore Papageorgiou (**2024**). Investment in Infrastructure and Trade: The Case of Ports. NBER Working Paper 32503, NBER Working Paper

Bu, Yisheng (**2006**). Fixed capital stock depreciation in developing countries: Some evidence from firm level data. *The Journal of Development Studies* 42(5), pp. 881–901. Place: Abingdon

Caselli, Francesco and James Feyrer (**2007**). The Marginal Product of Capital. *The Quarterly Journal of Economics* 122(2), pp. 535–568

Caselli, Francesco and Daniel J Wilson (**2004**). Importing technology. *Journal of Monetary Economics* 51(1), pp. 1–32

Caunedo, Julieta and Namrata Kala (**2023**). Mechanizing Agriculture. Working Paper

Caunedo, Julieta, Namrata Kala, and Haimeng Zhang (**2024**). Economies of Density and Congestion in Equipment Capital. Working Paper

Caunedo, Julieta and Elisa Keller (**2021**). Capital Obsolescence and Agricultural Productivity. *The Quarterly Journal of Economics* 136(1), pp. 505–561

Chamberlin, Edward H. (**1948**). Proportionality, Divisibility and Economies of Scale. *The Quarterly Journal of Economics* 62(2), pp. 229–262

Chari, Anusha and Jennifer Rhee (**2021**). The Return to Capital in Capital-Scarce Countries

Cramer, J. S. (**1958**). The Depreciation and Mortality of Motor-Cars. *Journal of the Royal Statistical Society* 121(1), p. 18

Dagsvik, John K. (**1994**). Discrete and Continuous Choice, Max-Stable Processes, and Independence from Irrelevant Attributes. *Econometrica* 62(5), pp. 1179–1205

de Mel, Suresh, David Mckenzie, and Christopher Woodruff (**2008**). Returns to Capital in Microenterprises: Evidence from a Field Experiment. *The Quarterly Journal of Economics* 123(4)

Diamond, Peter A. (**1982**). Aggregate Demand Management in Search Equilibrium. *Journal of Political Economy* 90(5), pp. 881–894

Donaldson, Dave and Richard Hornbeck (**2016**). Railroads and American Economic Growth: A "Market Access" Approach. *The Quarterly Journal of Economics* 131(2), pp. 799–858

Dupuy, Arnaud and Alfred Galichon (**2014**). Personality Traits and the Marriage Market. *Journal of Political Economy*

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer (**2015**). The Next Generation of the Penn World Table. *American Economic Review* 105(10), pp. 3150–3182

Feldstein, Martin S. and David K. Foot (**1971**). The Other Half of Gross Investment: Replacement and Modernization Expenditures. *The Review of Economics and Statistics* 53(1), pp. 49–58. Place: CAMBRIDGE

Feldstein, Martin S. and Michael Rothschild (**1974**). Towards an Economic Theory of Replacement Investment. *Econometrica* 42(3), pp. 393–423

Feller, William (**1957**). *An Introduction to Probability Theory and Its Application*. John Wiley, New York

Fraumeni, Barbara M (**1997**). The Measurement of Depreciation in the U.S. National Income and Product Accounts. Survey of Current Business, BEA

Gollin, Douglas and Christopher Udry (**2021**). Heterogeneity, Measurement Error, and Misallocation: Evidence from African Agriculture. *Journal of Political Economy* 129(1), pp. 1–80. Place: Chicago

Graff, Tilman (**2024**). Spatial inefficiencies in Africa's trade network. *Journal of Development Economics* 171, p. 103319

Grant, Matthew and Meredith Startz (**2020**). Cutting Out the Middleman: The Structure of Chains of Intermediation. Working Paper

Griliches, Zvi (**1962**). Capital Stock and Investment Functions: Some Problems of Concept and Measurement. In: Christ (Editor), *Measurement in Economics*. Stanford University Press, Stanford

Hall, Robert and Charles Jones (**1999**). Why do Some Countries Produce So Much More Output Per Worker than Others? *The Quarterly Journal of Economics* 114(1)

Henning, David and Joseph Okello (**2024**). Tax Audits and Their Distortionary Effects. Working Paper

Hirschman, Albert O (**1958**). *The Strategy of Economic Development*. Yale University Press, New Haven

Hsieh, Chang-Tai and Peter J. Klenow (**2007**). Relative Prices and Relative Prosperity. *American Economic Review* 97(3), pp. 562– 585

——— (**2009**). Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics* 124(4), pp. 1403–1448. Place: CARY

——— (**2010**). Development Accounting. *American Economic Journal: Macroeconomics* 2(1), pp. 207–223

Hulten, Charles R. and Frank C. Wykoff (**1996**). Issues in the Measurement of Economic Depreciation: Introductory Remarks. *Economic Inquiry* 34(1), pp. 10–23

Iowa State University Extension School (**2025**). Ag Decision Maker. `https://www.extension.iastate.edu/agdm/homepage.html`

Jones, Charles I (**2011**). Intermediate Goods and Weak Links in the Theory of Economic Development. *American Economic Journal: Macroeconomics* 3(2), pp. 1–28

Jorgenson, Dale W (**1963**). Capital Theory and Investment Behavior. *American Economic Review* 53(2)

Jorgenson, Dale W. (**1996**). Empirical Studies of Depreciation. *Economic Inquiry* 34(1), pp. 24–42

Kaboski, Joseph P, Molly Lipscomb, Virgiliu Midrigan, and Carolyn Pelnik (**2024**). How Important are Investment Indivisibilities for Development? Experimental Evidence from Uganda. Working Paper

Kaplan, Greg and Giovanni L. Violante (**2014**). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica* 82(4), pp. 1199–1239

Karabarbounis, Loukas and Brent Neiman (**2014**). Capital Depreciation and Labor Shares Around the World: Measurement and Implications. Technical Report w20606, National Bureau of Economic Research, Cambridge, MA

Kay, Ronald D., William M. Edwards, and Patricia Ann Duffy (**2016**). Farm management Publisher: [Singapore]: McGraw-Hill Education

Kleinman, Benny, Ernest Liu, and Stephen J Redding (**2023**). Dynamic Spatial General Equilibrium. *Econometrica* 91(2), p. 56

Klenow, Peter and Andres Rodriguez-Clare (**1997**). The Neoclassical Revival in Growth Economics: Has It Gone Too Far? In: *NBER macroeconomics annual 1997*. MIT Press, Cambridge, Mass. OCLC: 1421132709

Kremer, Michael (**1993**). The O-Ring Theory of Economic Development. *The Quarterly Journal of Economics* 108(3)

Kremer, Michael, Jean Lee, Jonathan Robinson, and Olga Rostapshova (**2013**). Behavioral Biases and Firm Behavior: Evidence from Kenyan Retail Shops. *American Economic Review P&P* 103(3), pp. 362–368

Lagos, Ricardo (**2000**). An Alternative Approach to Search Frictions. *Journal of Political Economy* 108(5)

Lanteri, Andrea (**2018**). The Market for Used Capital: Endogenous Irreversibility and Reallocation over the Business Cycle. *American Economic Review* 108(9), pp. 2383–2419

Lee, Kenneth, Edward Miguel, and Catherine Wolfram (**2020**). Experimental Evidence on the Economics of Rural Electrification. *Journal of Political Economy* 128(4), pp. 1523–1565

Lerner, Abba Ptachya (**1944**). *The Economics of Control: Principles of Welfare Economics*. The Macmillan Company, New York

Levhari, David and Eytan Sheshinski (**1970**). A Microeconomic Production Function. *Econometrica* 38(3), p. 559

Lewis, W. Arthur (**1955**). *The Theory of Economic Growth*. R.D. Irwin, Homewood, Ill.

Lucas, Robert E (**1990**). Why Doesn't Capital Flow from Rich to Poor Countries? *American Economic Review P&P* 80(2)

Maarek, Paul and Elsa Orgiazzi (**2020**). Development and the Labor Share. *The World Bank Economic Review* 34(1), pp. 232–257

Mankiw, David Romer, and David Weil (**1992**). A Contribution to the Empirics of Economic Growth. *The Quarterly Journal of Economics* 107(May), pp. 407–437

Marshall, Alfred (**1890**). *Principles of Economics*. Macmillan and Co., London, New York

McKenzie, David and Christopher Woodruff (**2008**). Experimental Evidence on Returns to Capital and Access to Finance in Mexico. *The World Bank Economic Review* 22(3), pp. 457–482

Melitz, Marc (**2003**). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica* 71(6), pp. 1695–1725

Miyauchi, Yuhei (**2024**). Matching and Agglomeration: Theory and Evidence From Japanese Firm-to-Firm Trade. *Econometrica* 92(6), pp. 1869–1905

Mohring, Herbert (**1972**). Optimization and Scale Economies in Urban Bus Transportation. *American Economic Review* 62(4)

Moll, Benjamin (**2014**). Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation? *American Economic Review* 104(10), pp. 3186–3221

Pelnik, Carolyn (**2024**). Moving to Profitability? Alleviating Constraints on Microentrepreneur Location. Working Paper

Raisaro, Claude (**2024**). Incentives Justifying Nonconformity: Experimental Evidence from Motortaxi Organizations in Uganda. Working Paper

Rampini, Adriano A. (**2019**). Financing Durable Assets. *American Economic Review* 109(2), pp. 664–701

Sattinger, Michael (**2002**). A Queuing Model of the Market for Access to Trading Partners. *International Economic Review* 43(2), pp. 533–547

Schündeln, Matthias (**2013**). Appreciating depreciation: physical capital depreciation in a developing country. *Empirical Economics* 44(3), pp. 1277–1290

Sen, Amartya Kumar (**1962**). On the Usefulness of Used Machines. *The Review of Economics and Statistics* 44(3)

Solow, R. M. (**1959**). Investment and Technical Progress. In: *Mathematical Methods in the Social Sciences*. Arrow, Karlin, Suppes

Solow, Robert M. (**1956**). A Contribution to the Theory of Economic Growth. *The Quarterly Journal of Economics* 70(1), pp. 65–94. ISBN: 1531-4650

Startz, Meredith (**2016**). The value of face-to-face: Search and contracting problems in Nigerian trade. Technical report

Thomas, Douglas S and Brian A Weiss (**2020**). Economics of manufacturing machinery maintenance: a survey and analysis of U.S. costs and benefits. Technical Report NIST AMS 100-34, National Institute of Standards and Technology, Gaithersburg, MD

Udry, Christopher and Santosh Anagol (**2006**). The Return to Capital in Ghana. *American Economic Review* 96(2), pp. 388–393. Place: NASHVILLE

United Nations (Editor) (**2009**). *System of national accounts 2008*. Number 2 in Studies in methods. United Nations [u.a.], New York, rev. 5 edition

Vitali, Anna (**2022**). Consumer Search and Firm Location: Theory and Evidence from the Garment Sector in Uganda. Working Paper

Walker, Michael, Nachiket Shah, Edward Miguel, Dennis Egger, Felix Samy Soliman, and Tilman Graff (**2024**). Slack and Economic Development. NBER Working Paper 33055

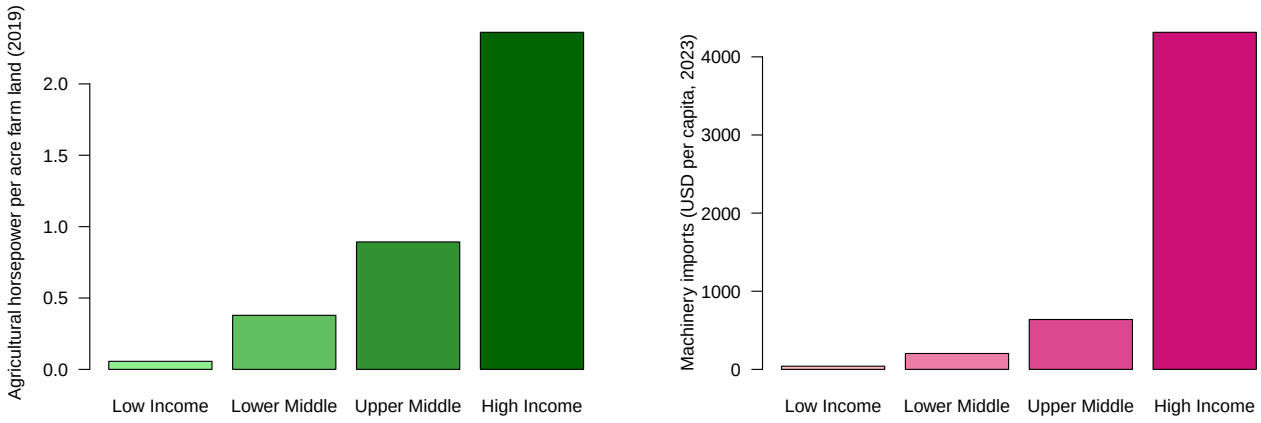
Weiss, Tim, Michael Lounsbury, and Garry Bruton (**2024**). Survivalist Organizing in Urban Poverty Contexts. *Organization Science* 35(5), pp. 1608–1640

Whitin, T. M. and M. H. Peston (**1954**). Random Variations, Risk, and Returns to Scale. *The Quarterly Journal of Economics* 68(4), p. 603

World Bank (**2019**). World Bank Group's ENTERPRISE SURVEY, Understanding the Questionnaire. `https://www.enterprisesurveys.org/content/dam/enterprisesurveys/documents/methodology/ES_QuestionnaireManual_2019.pdf`

A Additional figures

Figure A.1: Rich countries use more capital than poor countries

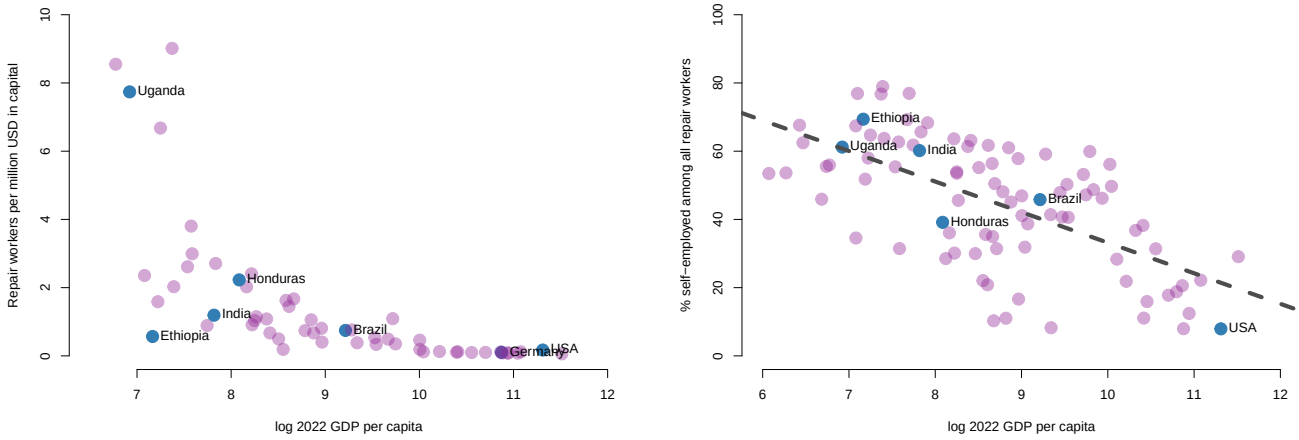


(a) Agricultural horsepower per acre farm land

(b) Machinery imports per capita

Notes: Trends of global differences in capital density between rich and poor countries. Data on agricultural horsepower from Our World in Data, data on machinery imports from WIOTs.

Figure A.2: Repair workers in poor countries are more numerous and self-employed

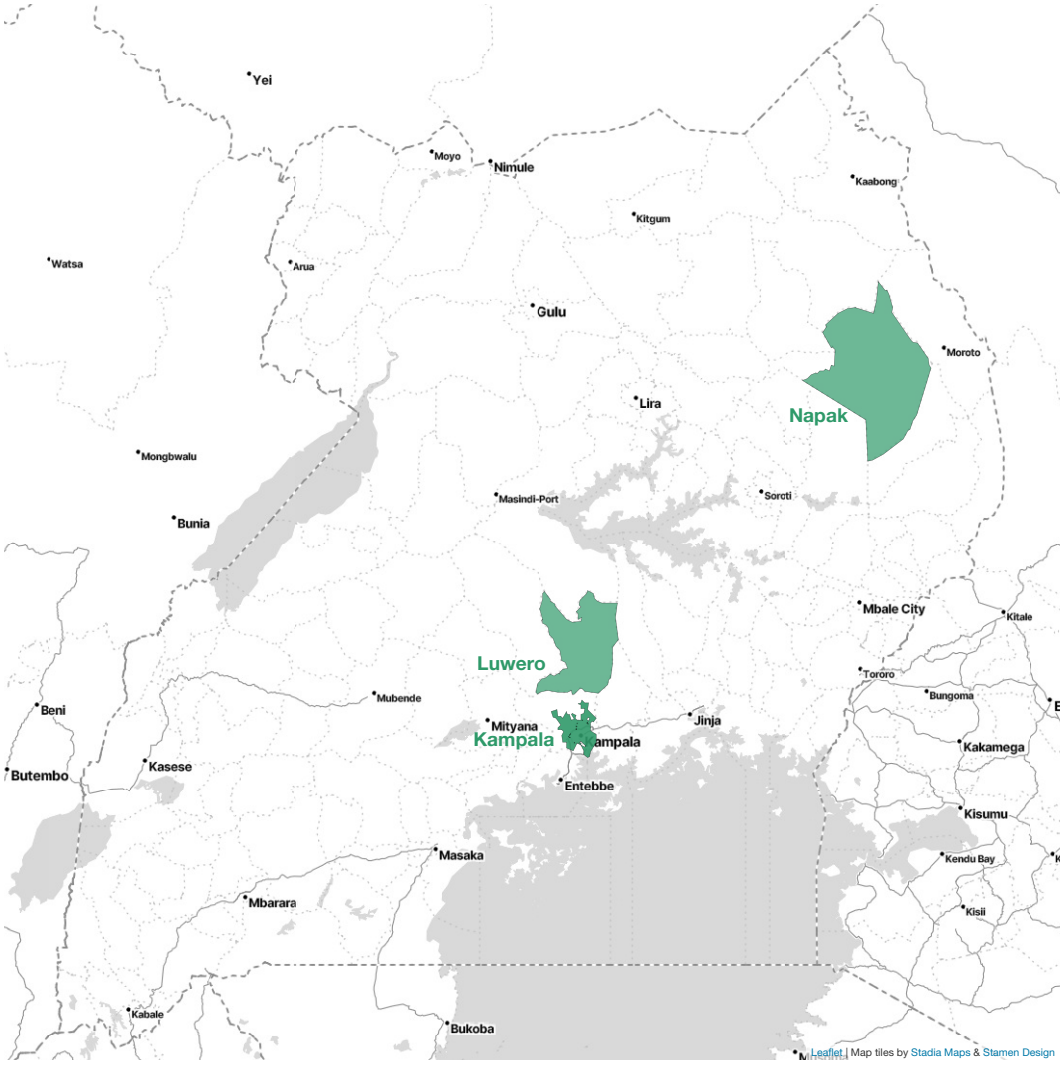


(a) Number of mechanics per $1 million in machinery

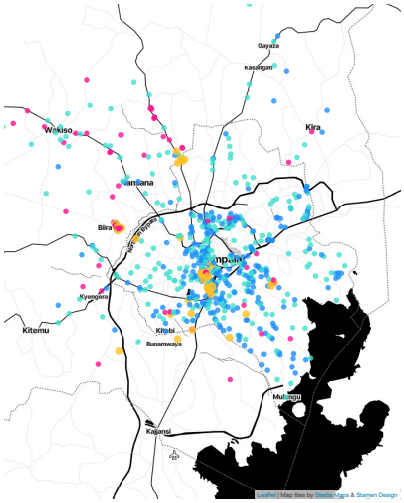
(b) Share of mechanics who are self-employed

Notes: Global patterns on the size and employment distribution of the repair sector. Panel a prints the number of workers in mechanic and repair occupations (ISIC codes 33, 452, and 95) divided by a country's total capital stock (as reported by the Penn World Table, Feenstra et al., 2015). Magnitudes are such that Uganda, for example, has eight repair workers per one million USD worth of capital. Note that the data on capital stocks in the Penn World Table is constructed using a constant depreciation rate, an assumption I relax in this paper. Panel b prints the share of all repair workers who are self-employed (as opposed to being employees in a larger firm).

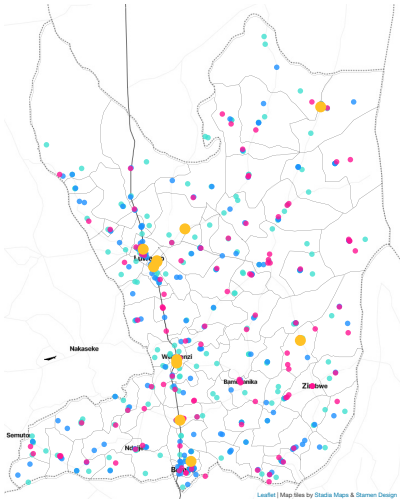
Figure A.3: Detailed survey locations



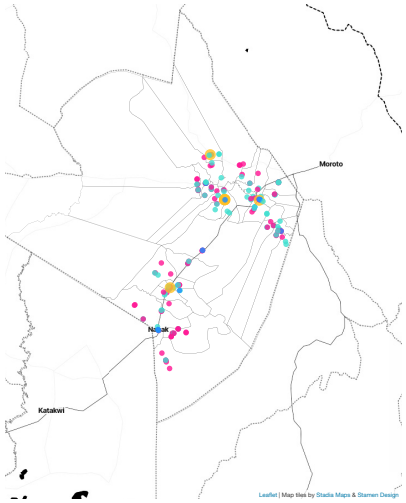
(a) Study locations within Uganda



(b) Kampala + suburbs



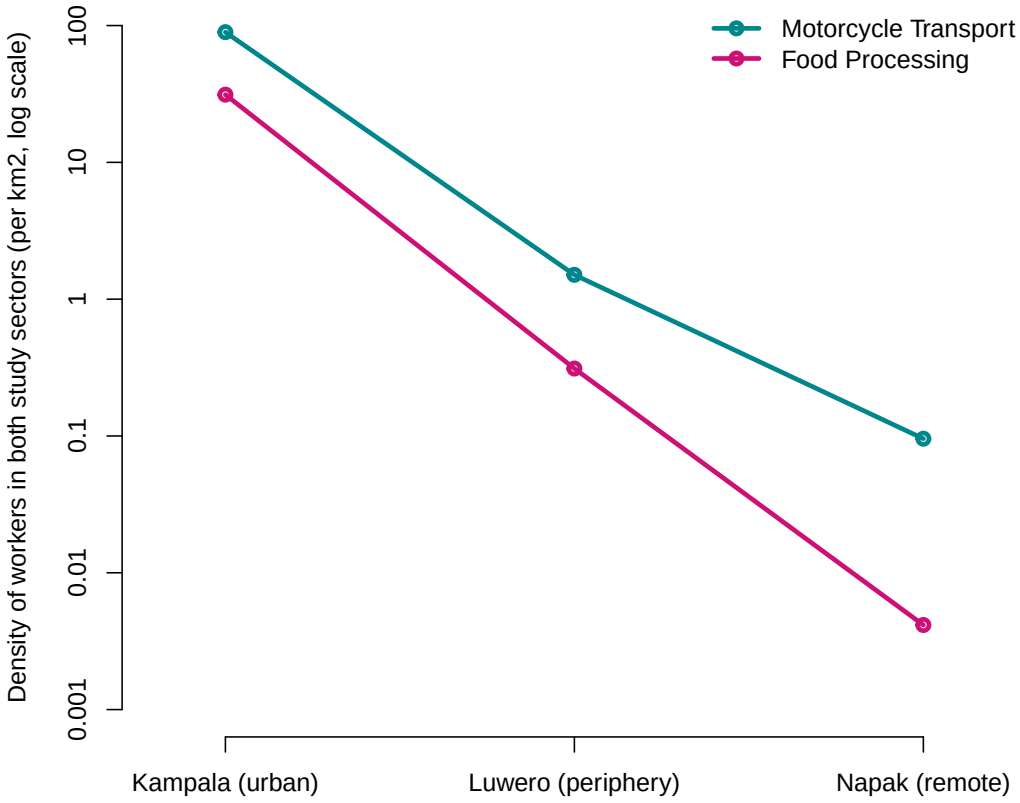
(c) Luwero district



(d) Napak district

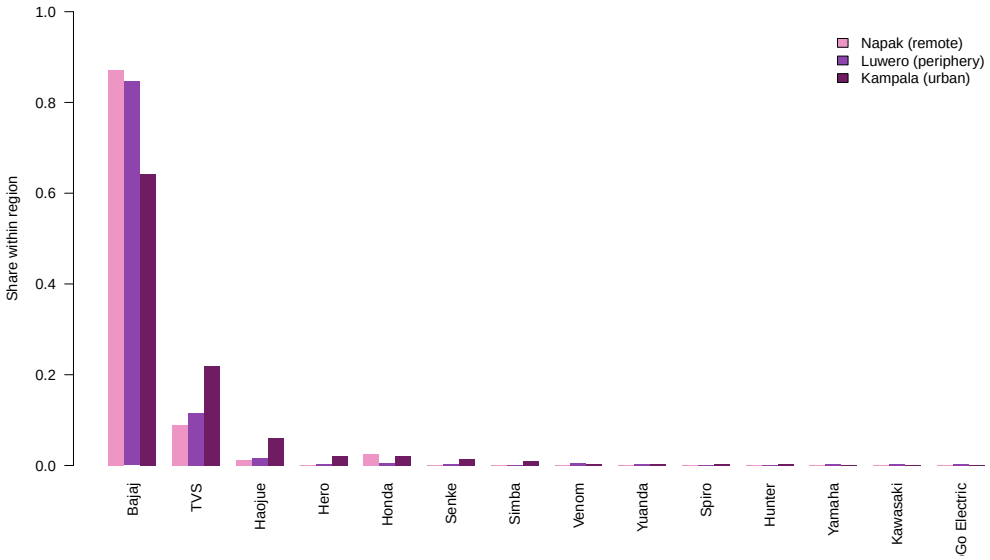
Notes: Detailed maps of survey locations in all three regions. Colors represent motorcycle riders (green), food processing firms (pink), motorcycle mechanics (blue), and food processing mechanics (yellow).

Figure A.4: Density of workers across study sectors and regions

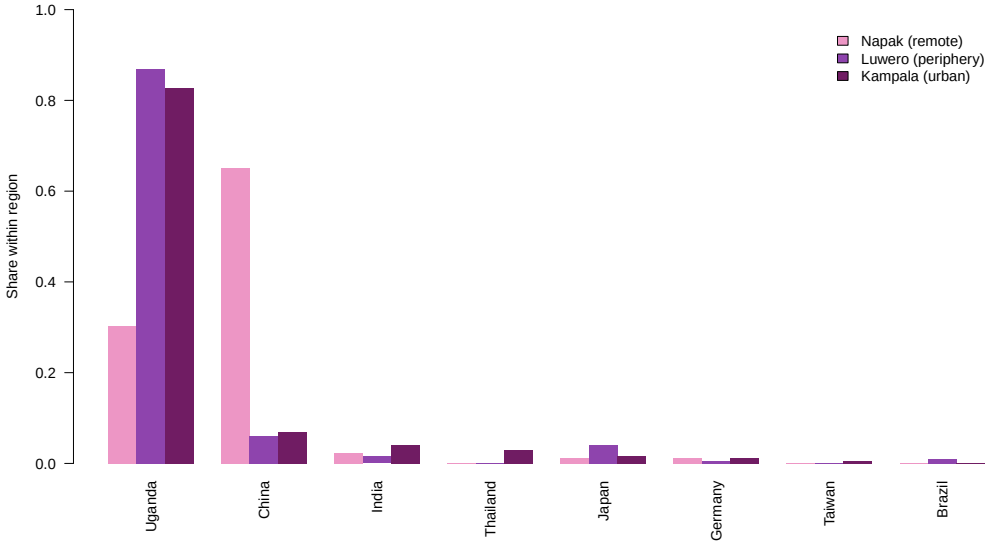


Notes: Number of workers in motorcycle transport and food processing per square kilometer across the three study districts. Data from the 2014 National Population and Housing Census. "Food processing" in the Census is a broader category than mine and so might include animal feed processors and other related trades.

Figure A.5: Distribution of motorcycle brands and food processing machines across regions



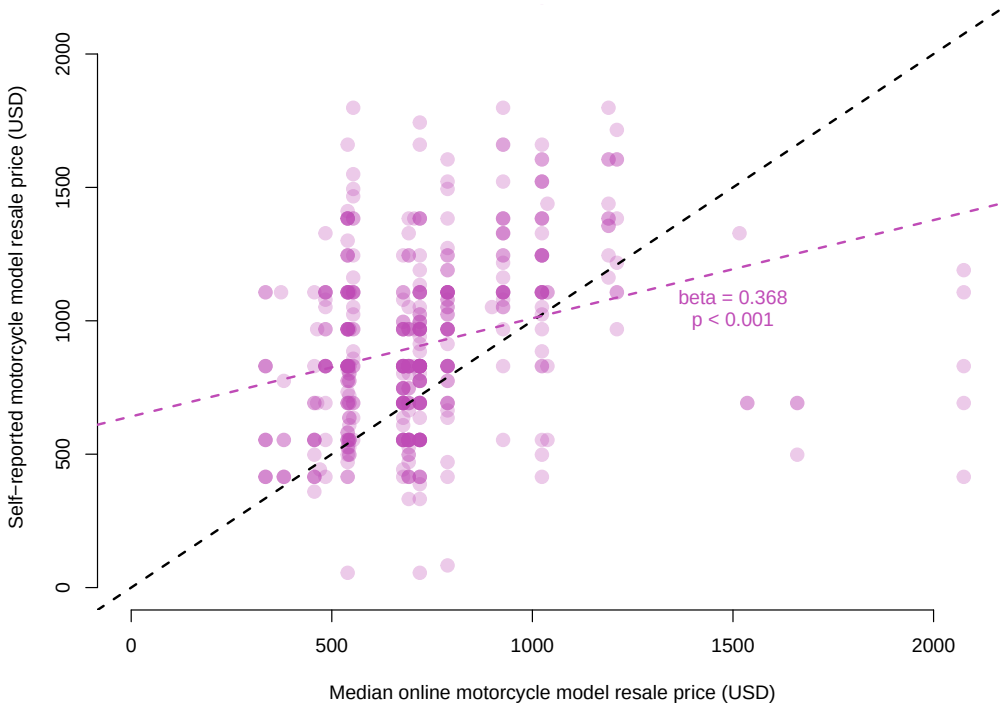
(a) Distribution of motorcycle brands by region



(b) Distribution of food processing machine countries of origin by region

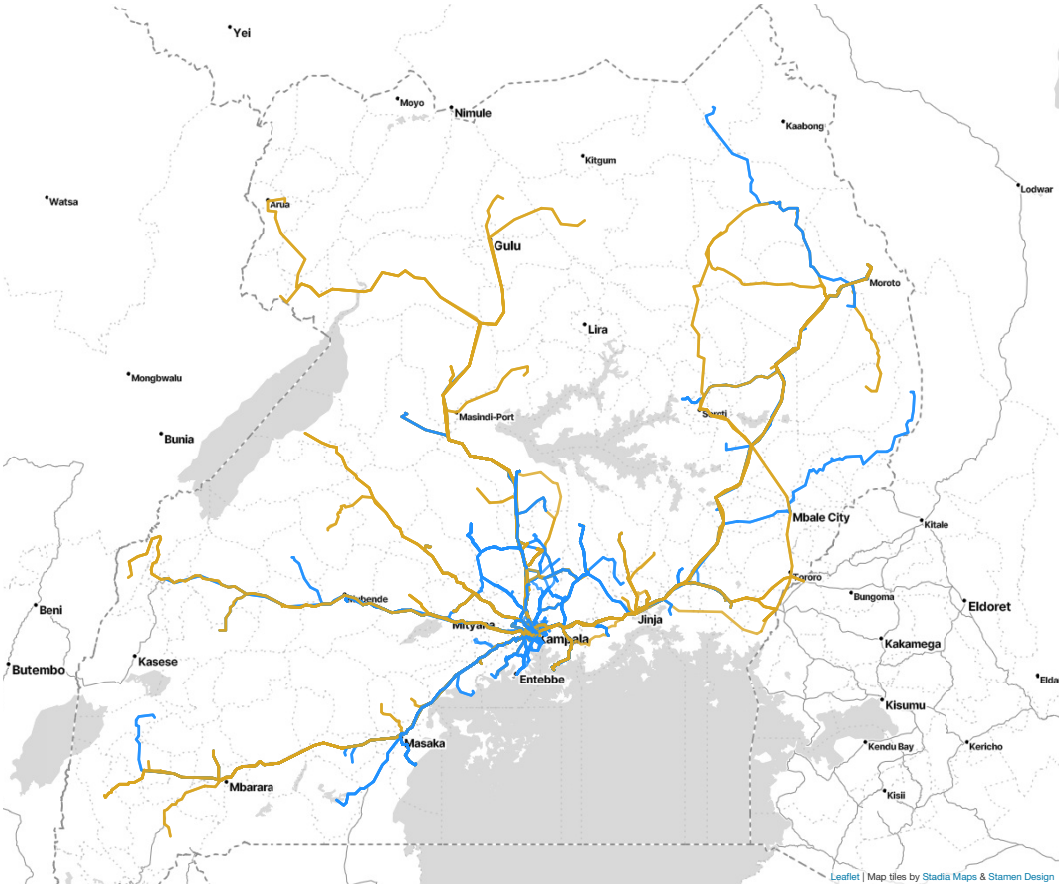
This figure prints the full distribution of motorcycle makes and food processing machine production countries by survey area. Among motorcycles, Bajaj is the most common brand but the market is substantially less concentrated in less remote areas. Among food processing machines, Ugandan-made machines dominate in two out of the three study regions, with Chinese made brands the most common in Napak district.

Figure A.6: Motorcycle riders' estimated machine values largely reflect observable resale price.



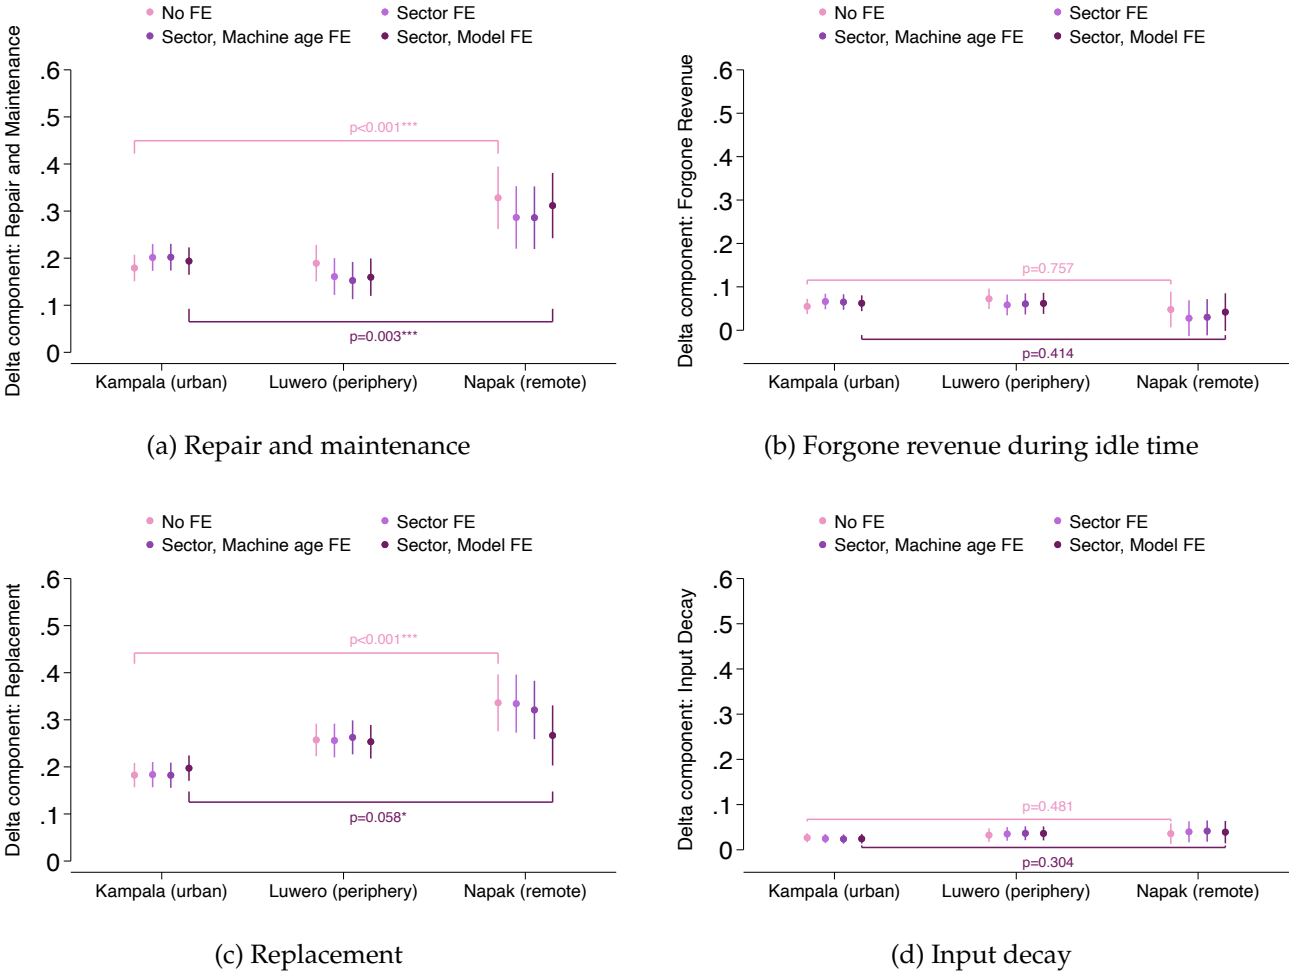
This figure prints motorcycle riders' self-reported estimates of the resale value of their motorbike against a sample of observed used motorcycle prices from a large Ugandan online marketplace. 45-degree line in black, line of best fit, achieved through projection of $\text{OnlinePrice}_b = \beta_0 + \beta\text{EstimatedPrice}_{i,b} + \epsilon_{i,b}$ for machine i in age-make bin b, in purple. Section C.3 provides more details.

Figure A.7: Mechanics travel patterns.



Notes: Map of the reported travel patterns of motorcycle mechanics (blue) and food processing mechanics (yellow) in the last year.

Figure A.8: Components of Δ across regions



(a) Repair and maintenance

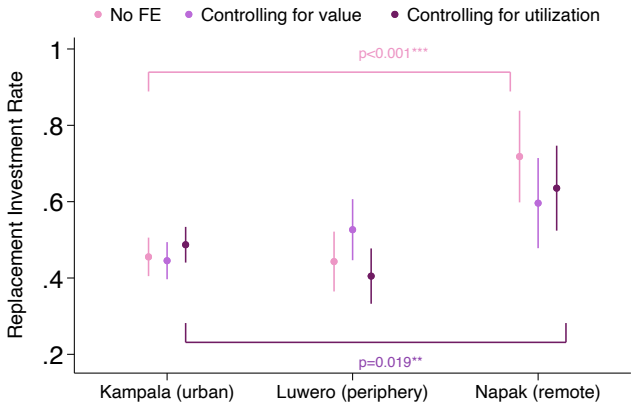
(b) Forgone revenue during idle time

(c) Replacement

(d) Input decay

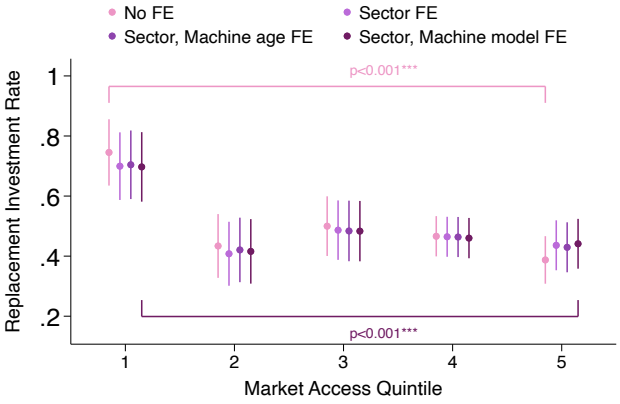
Notes: The four components of the replacement investment rate Δ split by survey region.

Figure A.9: Δ controlling for value and utilization



Notes: Robustness analyses displaying the spatial gradient in Δ after controlling for machine resale value (middle set of estimates) and machine utilization (rightmost set). The gradient persists, yet is weakened after controlling for these constituent components of Δ.

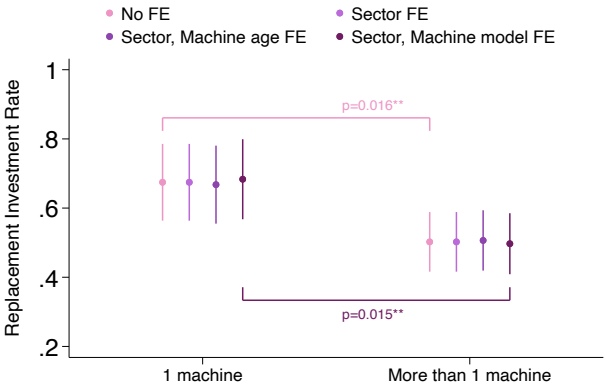
Figure A.10: Further evidence on Δ



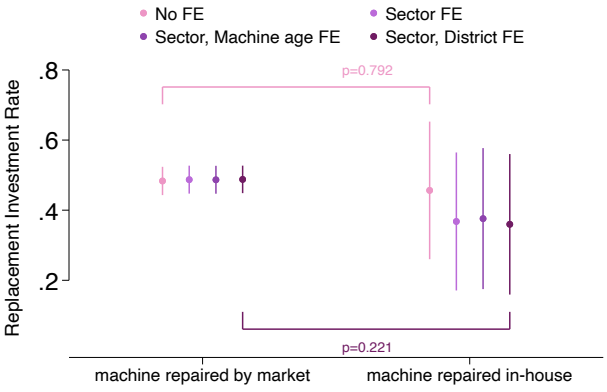
(a) Δ by market access quantiles



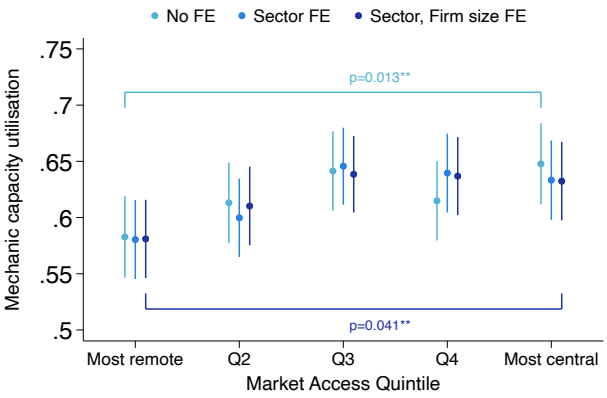
(b) Δ by number of employees



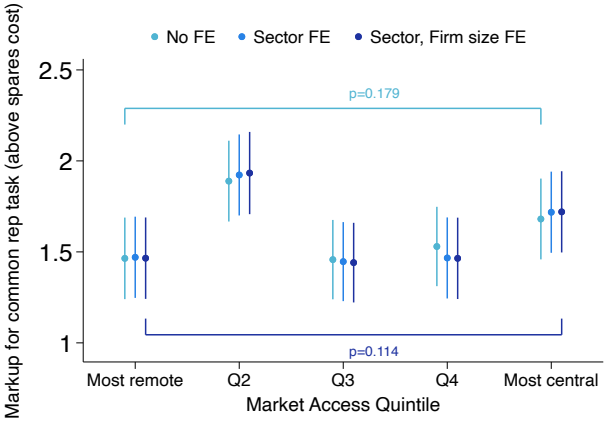
(c) Δ by number of machines



(d) Δ by whether machine repaired via market



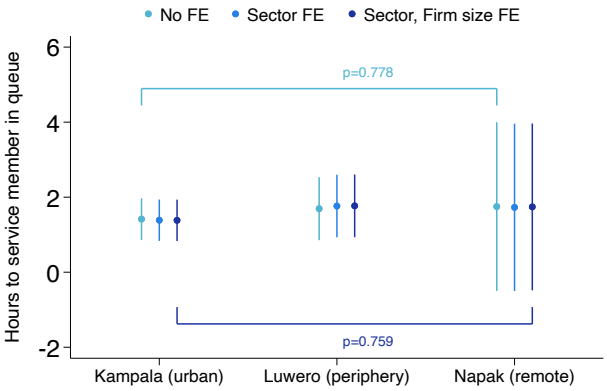
(e) Mechanic capacity utilization



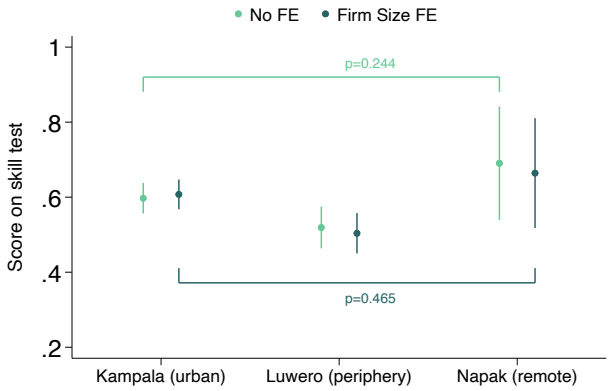
(f) Mechanic markup over spare parts cost

Notes: Additional evidence on how Δ varies across space and the firm distribution. Market access in panel a is computed according to Donaldson and Hornbeck (2016), which is further discussed in appendix C.5. In panel d, machine is considered as repaired in-house if either the firm has a mechanic on staff who is primarily responsible for repairing machines once they break down (around 5% of the firms in the sample) or if the respondent was able to repair the machine on their own when it last broke.

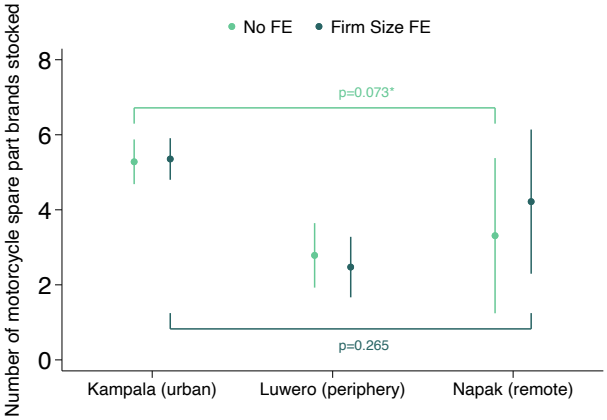
Figure A.11: Evidence on mechanic behavior by survey district



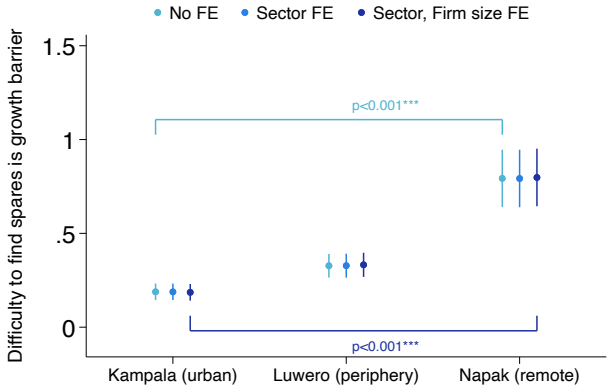
(a) Time to serve arriving customers



(b) Motor mechanics skills test results



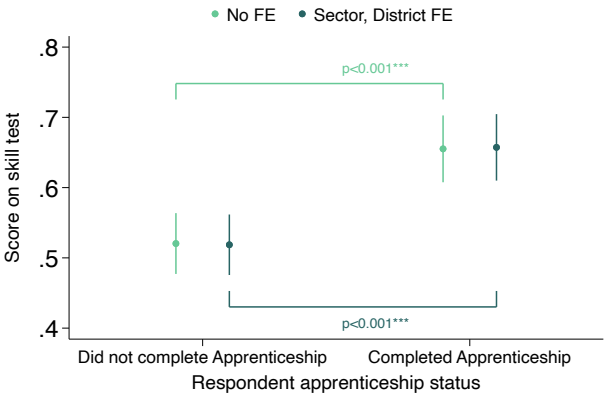
(c) Brands of spare parts in store



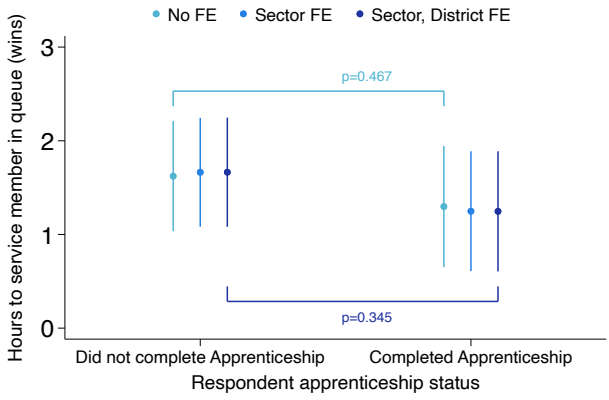
(d) Sourcing spare parts cited as problem

Notes: Identical analyses as in Figure 2, except split by the study district, not the market access quintile. As there are very few mechanics in Napak, confidence intervals on the leftmost set of estimates are large.

Figure A.12: Evidence on mechanic behavior by whether respondent completed an apprenticeship



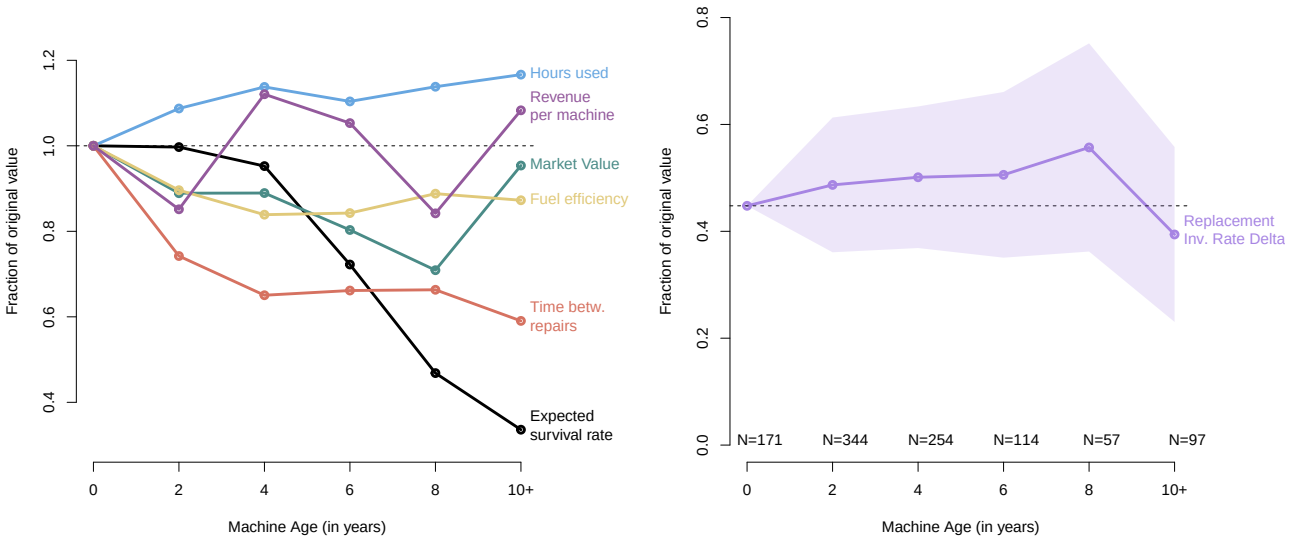
(a) Motor mechanics skills test results



(b) Time to serve arriving customers

Notes: Same analyses as panels 2b and 2a, except split by whether the mechanic completed an apprenticeship as part of their education (which is around half of mechanics).

Figure A.13: Machine age profiles

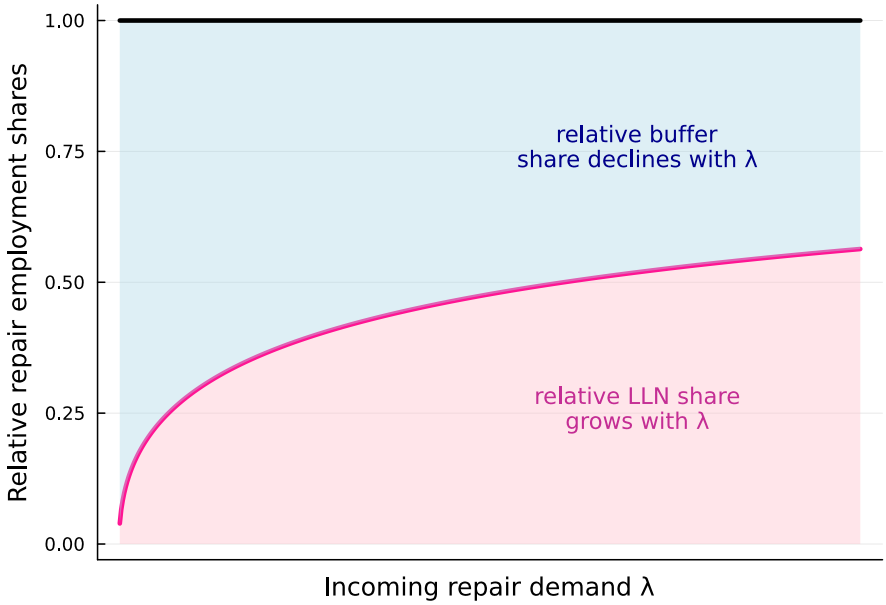


(a) Machine characteristics

(b) Replacement investment rate Δ

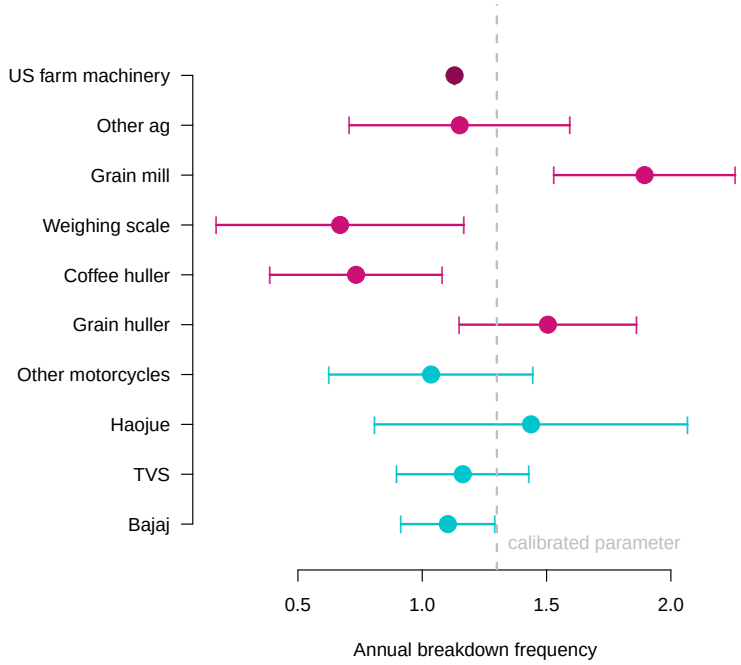
Notes: Observational age profiles of various machine characteristics and the replacement investment rate Δ. All gradients normalize by machine type (ie. motorcycle model and food processing machine class). Gradients obtained by running a regression of machine-level outcome Y on a series of 2-year age bin fixed effects. Machines aged 0-2 years identified through the intercept. Panel a normalizes every outcome by the intercept to compare relative changes. All outcomes are printed so that values above 1 indicate improvements in the outcome. For example, the yellow line implies machines at 10+ years are about 15% less fuel efficient than those 0-2 years old. Panel b indicates at the bottom the number of observations within each age bin, "10+" groups together all machines older than 10 years (and hence includes machines up to 23 years old), as well as the 95% confidence interval around the Δ age profile.

Figure A.14: Relative shares of buffer and predictable repair capacity

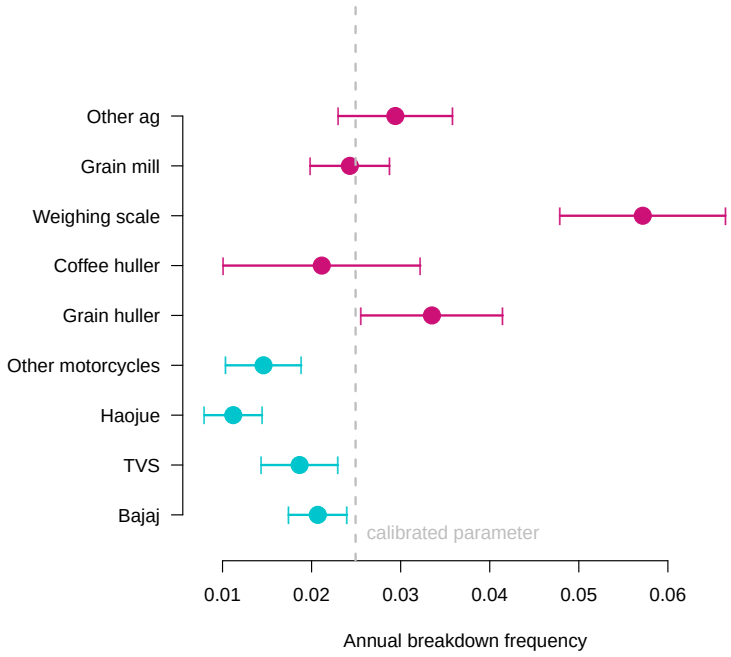


Notes: Equivalent to Figure 4a, yet expressing the components of the optimal repair employment by the minimum viable employment λ/μ (in pink), and the buffer employment in blue. Because buffer capacity grows with the square root of demand $\sqrt{\frac{\lambda r}{\mu\omega}}$, it is relatively more prominent at low levels of demand. Servicing an exceedingly rare event is hence more expensive.

Figure A.15: Model parameters by machine type



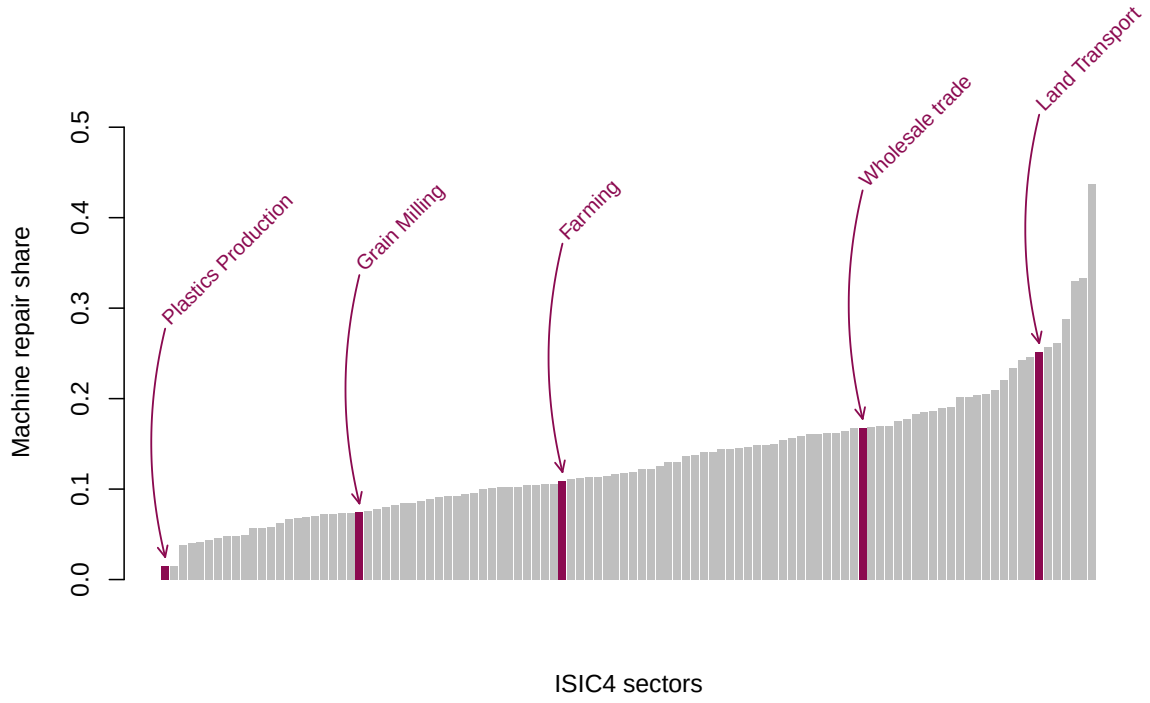
(a) Breakdown frequency δ_0 by machine type



(b) Relative spare parts price p_S/p_K by machine type

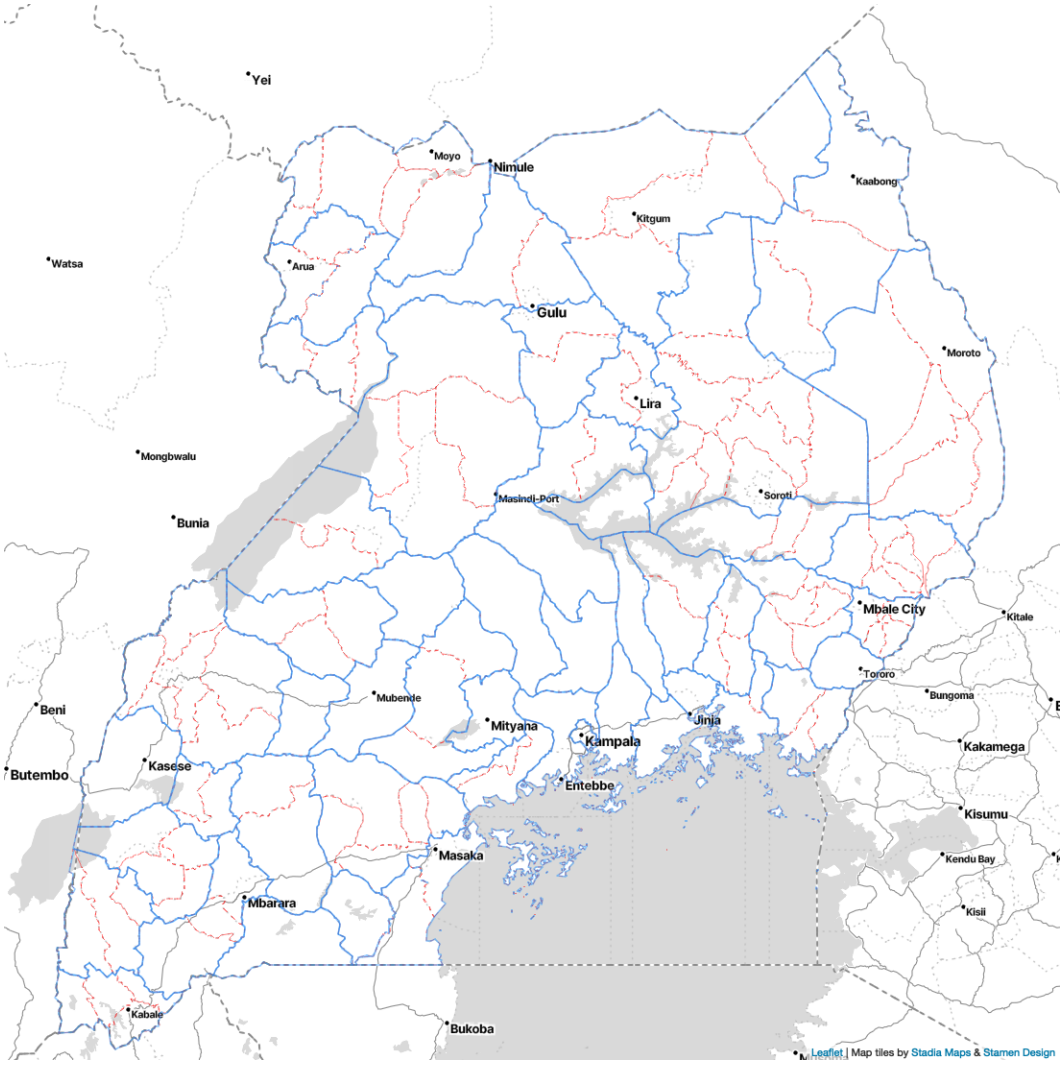
Notes: Machine-type level estimates of the rate of annual breakdown δ_0 and the relative costs of spare parts as a fraction of machine costs p_S/p_K. Panel a also prints the US estimate derived from ASABE Standards (2011).

Figure A.16: Relative machinery repair expenditures from tax data across Ugandan industries



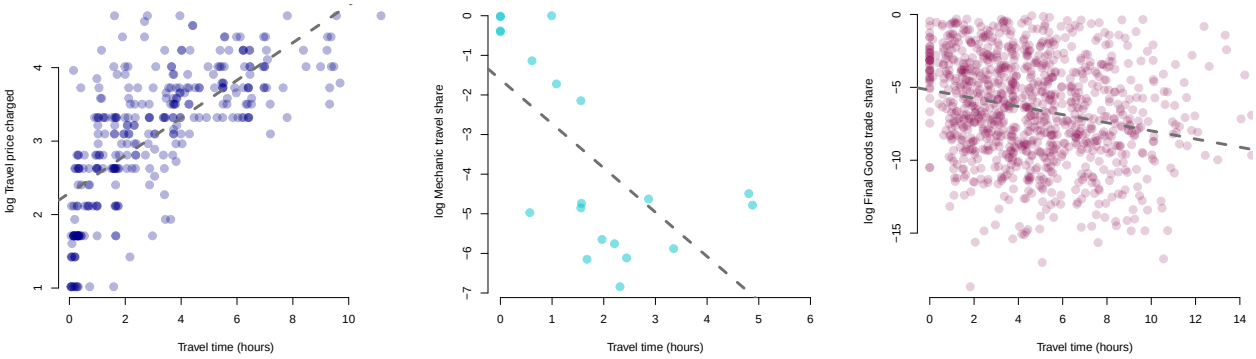
Notes: Distribution of formal-sector machinery repair expenditures normalised by total machinery stock value across ISIC-4 sectors. Data from administrative tax returns provided by the Ugandan Revenue Authority (URA), grouped over all tax years 2014–2023. Machinery repair share is constructed as the total expenditure on machine repair (as indicated on firms' corporate income tax return), divided by the total book value of machinery (from firms' balance sheets). The component of Δ accruing to repair expenditures in the main survey is 11% for motorcycles (compared to 25% for formal firms in land transport here) and 25% for grain milling (compared to 8% here).

Figure A.17: Model geography.



Notes: In order to fit the model to the Ugandan geography, I use a coarsened set of administrative locations. I transform the 138 districts of Uganda into 51 locations by combining adjacent areas. In this map, red district borders are eliminated to join districts together into larger units, described by the blue borders.

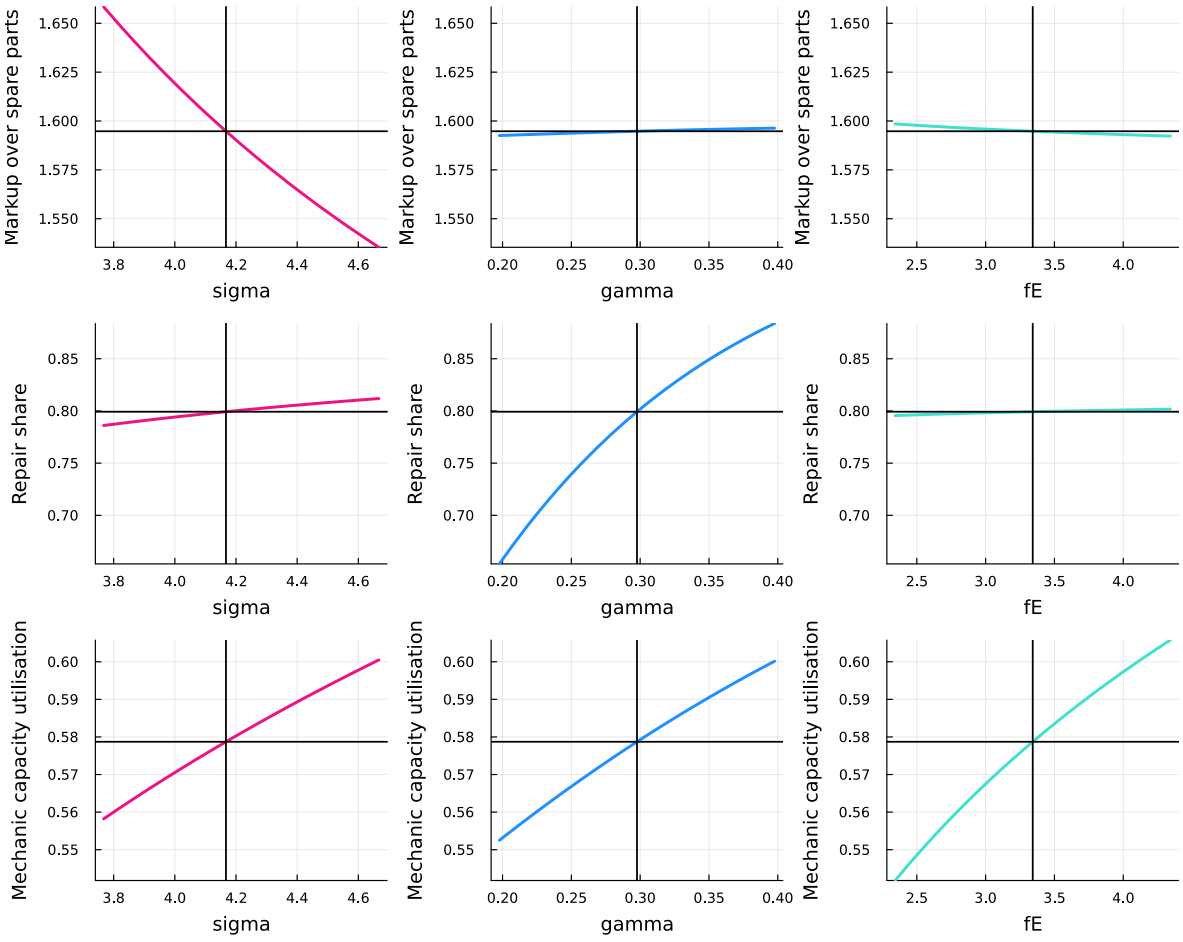
Figure A.18: Gravity regressions to isolate κ_M, ξ, and κ_G



(a) Mechanic price (b) Mechanic travel gravity (c) Final goods gravity

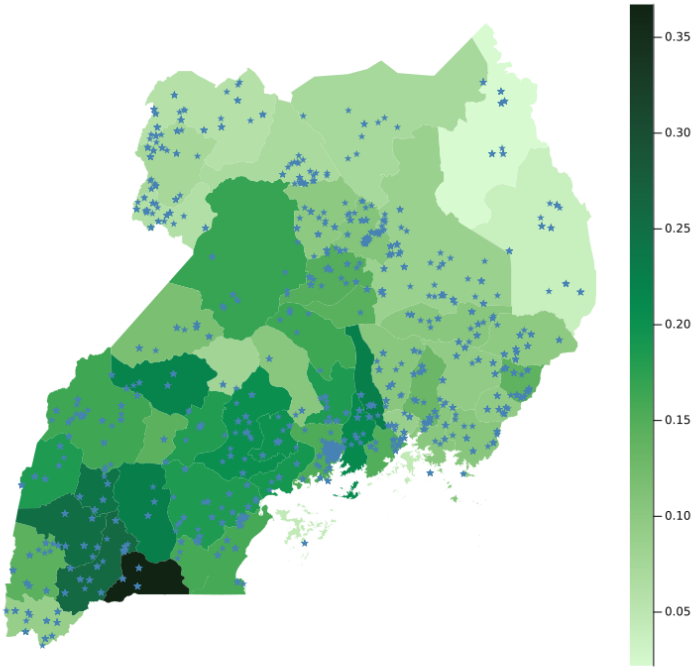
Notes: Scatter plots visualizing the structural gravity regressions run to estimate κ_M, ξ, and κ_G. The travel duration in hours between two districts dur_{ij} is on the x-axis of all three panels. Panel a prints $\log t_{ij}$ used to estimate equation 13, panel b prints $\log \Xi_{ij}$ for equation (14), and panel c prints $\log \Theta_{ij}$ used for equation (15).

Figure A.19: Local pseudo-"Jacobians" for Simulated Minimum Distance

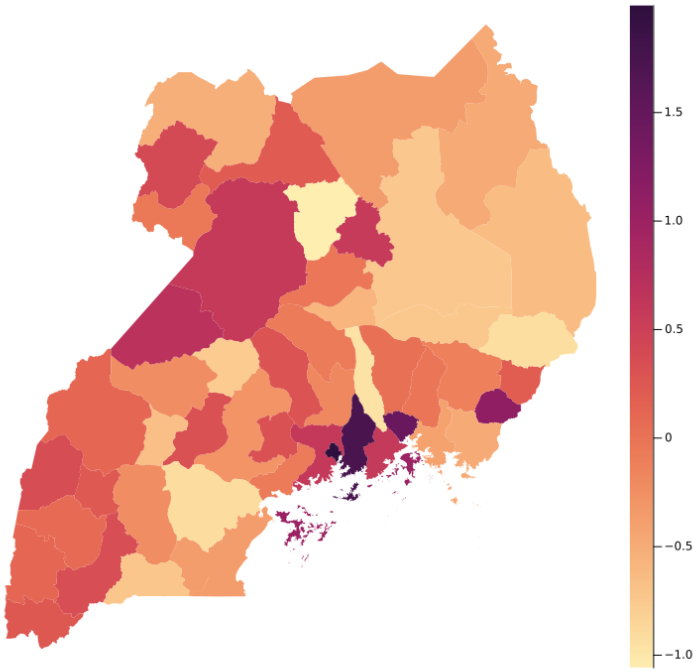


Notes: This figure documents how the three targeted moments are influenced by the three structural parameters. The markup mechanics charge over spare parts (first row) is mainly influenced by the mechanic elasticity of substitution σ: high σ means less market power for mechanics, who charge lower markups. The share of all breakages resulting in a repair (second row) is higher if the repair-replace parameter γ is higher and the likelihood of drawing a large value of the issue distribution ι is lower. Lastly, the average capacity utilization of mechanics (defined as the share of time a mechanic firm is actively working on a client task, third row) is influenced by all three parameters, including the mechanic entry fixed cost f_E. Higher f_E means mechanics only enter if they expect a large amount of orders and little idle time, driving up equilibrium utilization rates. In joint estimation, the vector (σ, γ, f_E) is chosen to match all three moments in the data, indicated by the horizontal lines.

Figure A.20: Spatial distribution of model-implied economic fundamentals



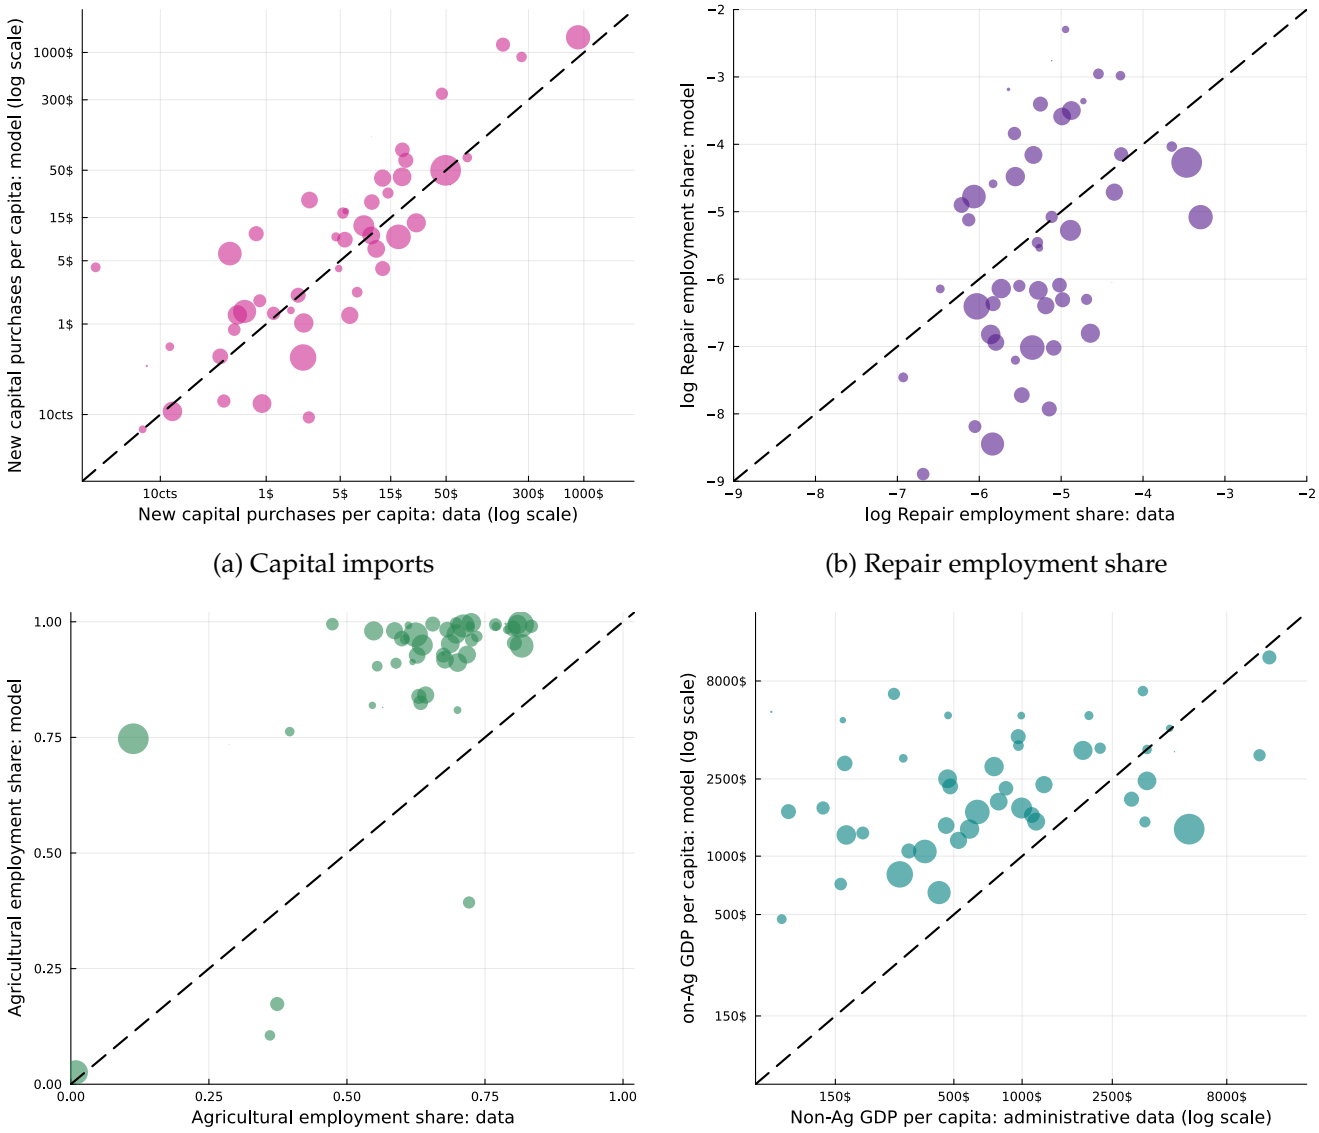
(a) Agricultural TFP Λ_ℓ (LSMS-ISA survey locations as dots)



(b) Inverted final goods productivity log A_ℓ

Notes: Spatial distribution of agricultural TFP Λ_ℓ and final goods TFP A_ℓ across Ugandan districts. Panel a prints agricultural TFP calculated from the World Bank Living Standards Measurement Survey – Integrated Survey on Agriculture (LSMS-ISA) module and using the agricultural production function Gollin and Udry (2021) estimate on the same data. Survey locations are printed as blue dots. Panel b prints inverted measures of (log) final goods TFP computed using the procedure outlined in appendix section E.1.

Figure A.21: Model fit



(a) Capital imports

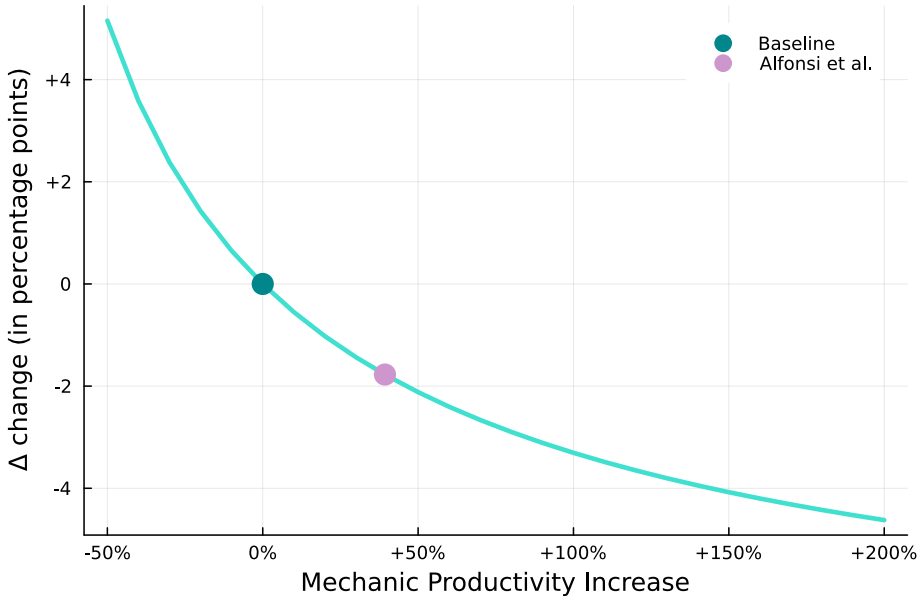
(b) Repair employment share

(c) Agricultural employment share

(d) Non-agricultural formal sector output per capita

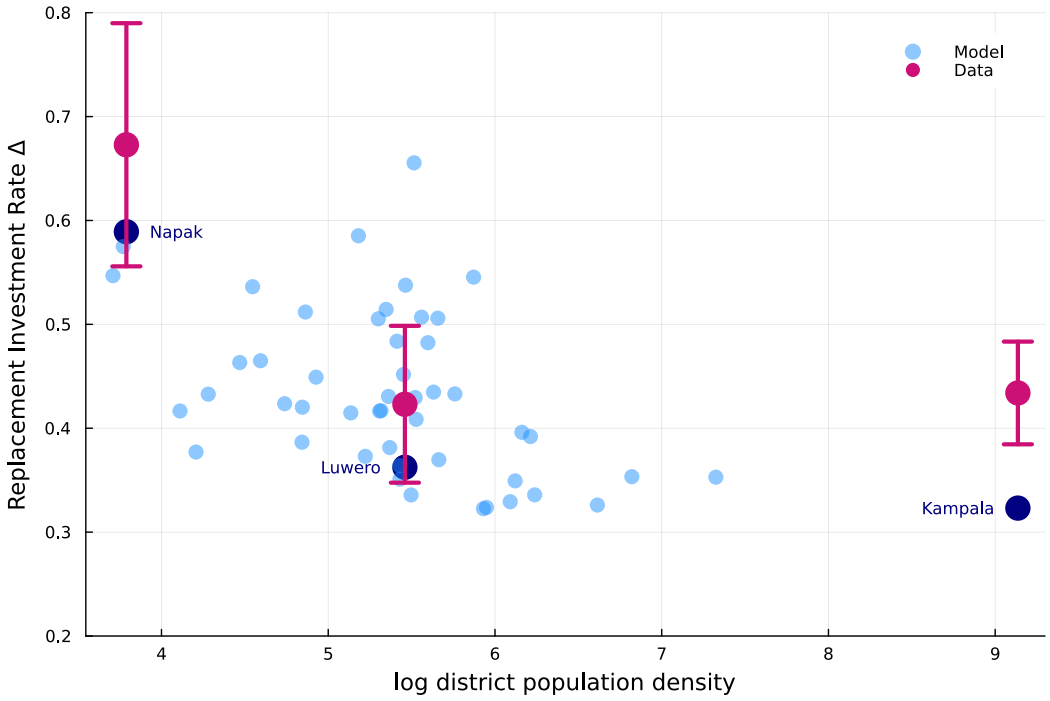
Notes: Fit of the calibrated model across Ugandan regions. Each dot represents a Ugandan region, scaled by the log of its population. Panel a compares model predictions of total capital imports (per capita, on a log scale) against administrative customs records of total machinery imports into each region from the Ugandan Revenue Authority. I rescale model predictions down by 44%, which is the share of machines in my sample that are Ugandan made and hence not imported. Panels b and c print employment shares in repair occupations and subsistence agriculture, derived from the 2014 Census. Panel d constructs a measure of per capita GDP in each region from administrative tax returns. I compile the total sales and the total wage bill of formal firms in each region. For sales, I subtract imported intermediates from other regions as reported in VAT records. For wages, I scale up by $1/(1-0.623) = 2.65$ to account for informal employment, which is estimated at around 63% in Uganda. I then take the maximum of both of these estimates to fill data gaps and issues in coverage of the administrative data especially in remote regions.

Figure A.22: Change in Δ in response to uniform increases in mechanic productivity μ



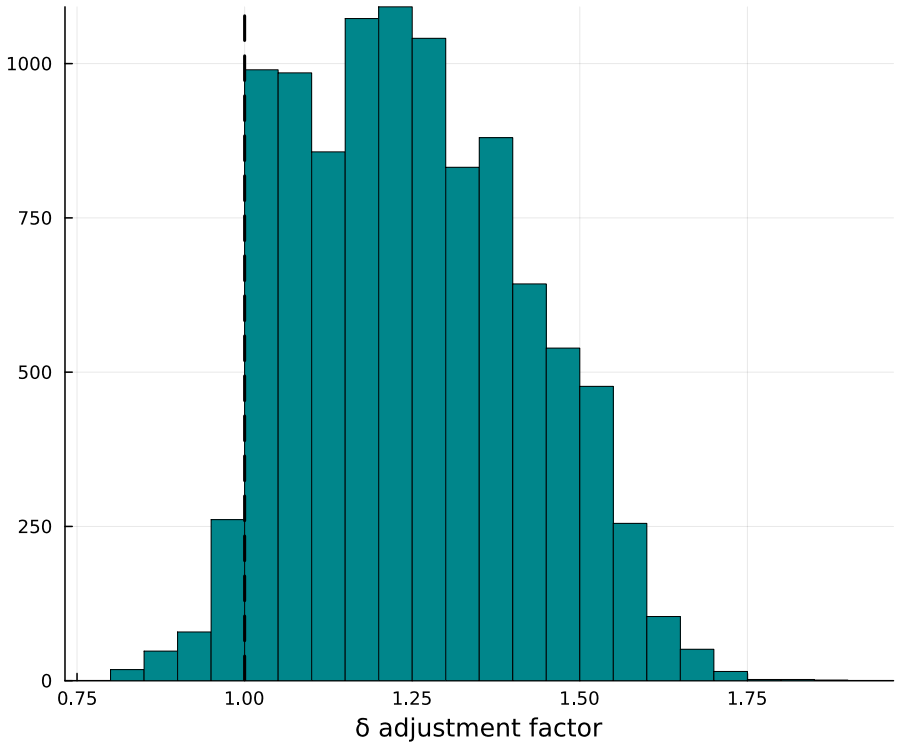
Notes: Counterfactual simulations of the resulting replacement investment rate if mechanic productivity μ was changed by various amounts. To visualize the Baumol (1967)-effect, note how Δ does not drop much further even if μ gets increased by plus 200%. On the other hand, the bottleneck status of repairs becomes apparent if μ drops even further and Δ shoots up on the left side of the graph. The pink dot corresponds to the effect presented in Figure 7, with μ increased by 39%.

Figure A.23: Δ against population density, calibrated with constant μ



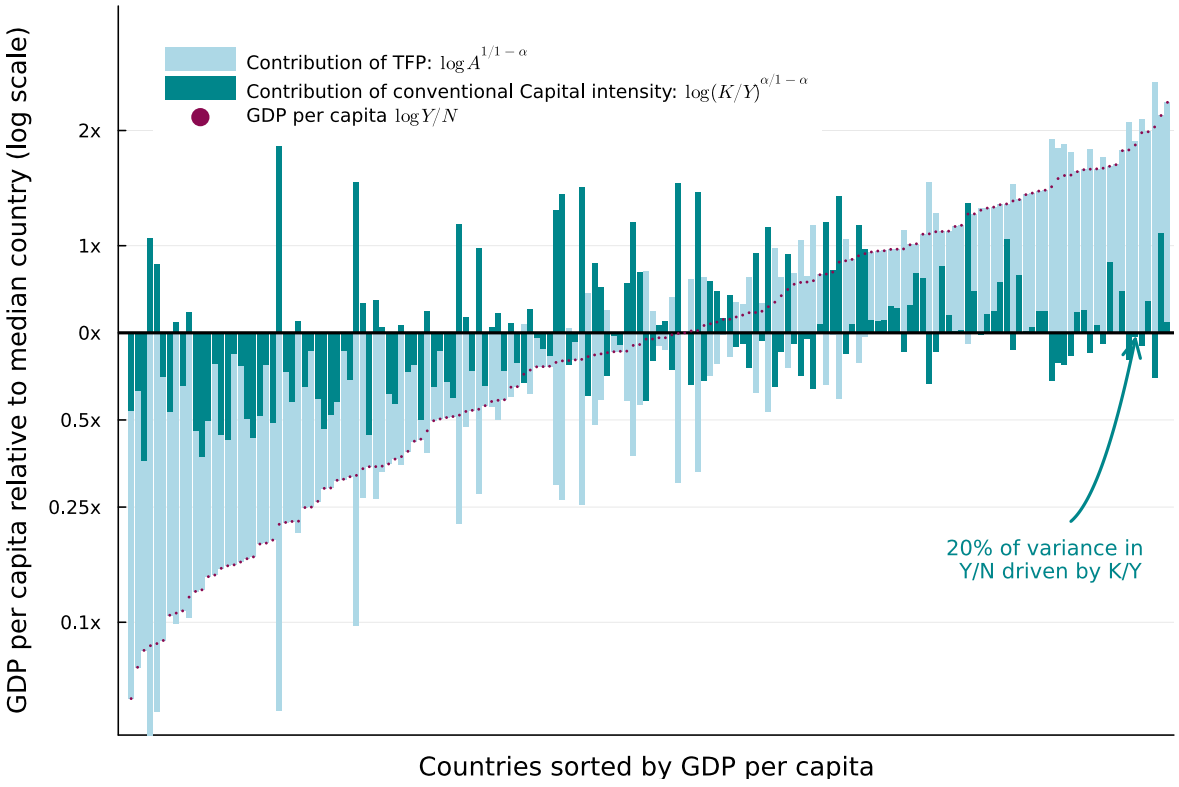
Notes: Equivalent to Figure 5a in the main text, plotting model-implied replacement investment rates Δ_ℓ against district population density. In the main model calibration, repair personnel productivity μ_ℓ is a linear mapping of local TFP A_ℓ (while ensuring that the average μ reflects the survey evidence of 194 repairs / year). As discussed in section 6.3.2, this extension instead keeps μ fixed at $\mu = 194$ everywhere. The gradient in Δ with remoteness is slightly weaker than in the full calibration, yet nevertheless pronounced.

Figure A.24: δ-adjustment factors across space and time

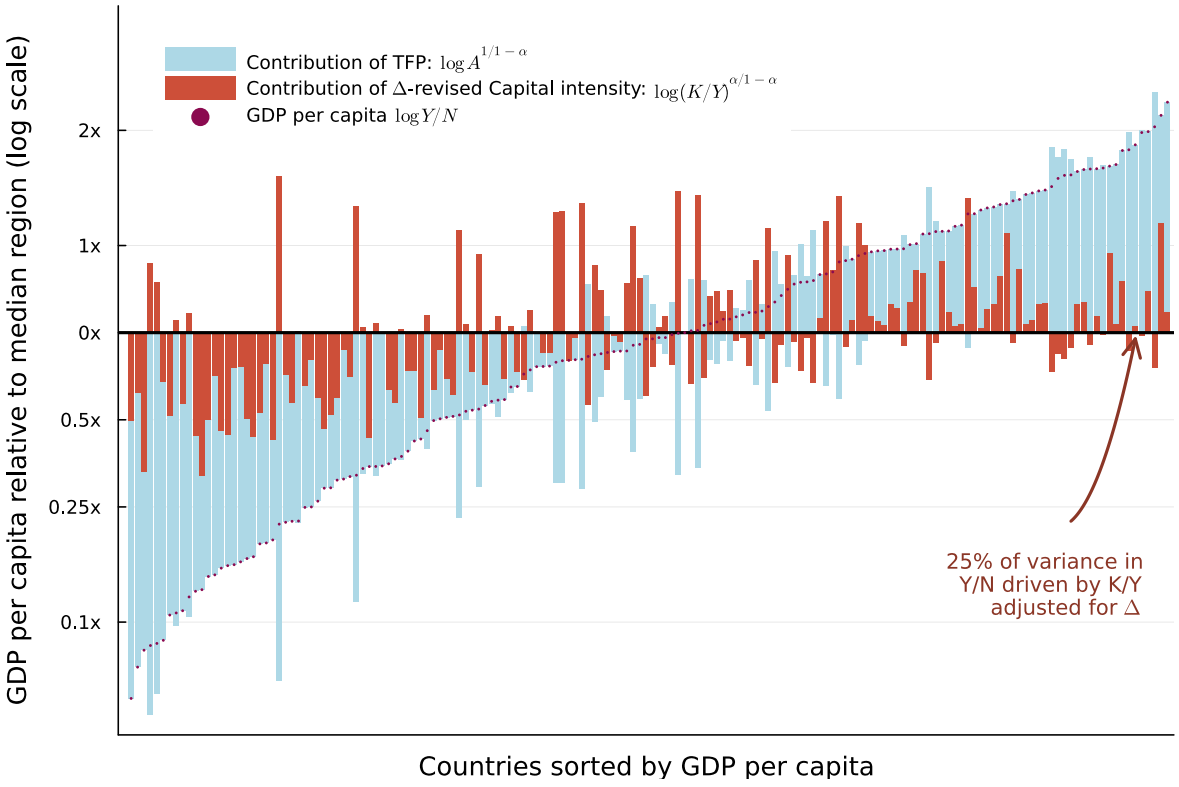


Notes: Histogram of the year-by-country specific adjustment factors of the depreciation rate. Factors above 1 imply the depreciation rate is increased.

Figure A.25: Global development accounting using conventional and updated capital stocks



(a) Contributions of K/Y and TFP to Y/N using conventional capital stocks



(b) Contributions of K/Y and TFP to Y/N using revised capital stocks

Notes: Figure identical to 6, but decomposing income differences across global countries instead of Ugandan regions. Panel a uses the conventional K/Y series from the Penn World Table. Panel b uses my revised series, which accounts for more of the differences in global income levels.

B Tables

Table A.1: Summary statistics by sector

	Motorcycle Riders	Food Processing	Motor Mechanics	Food Mechanics
Firm employment	1.04	2.71	2.42	2.16
Firm age (years)	4.45	6.29	5.26	8.97
Monthly firm revenue (USD)	207	1293	459	412
Monthly firm costs (USD)	99	829	228	196
Daily customers	13	10	15	2.14
Capacity utilisation	0.63	0.56	0.64	0.51
Revenue ratio best/worst week	2.51	3.33	2.98	3.87
Number of machines	1.00	1.61		
Machine purchase price (USD)	1241	1961		
Machine is imported from abroad	0.93	0.30		
Machine age (years)	3.64	5.94		
Machine hours in use yesterday	9.56	6.11		
Machine owned by firm	0.96	1.00		
Machine minutes rented out yesterday	8.05	0.00		
Has a mechanic on staff	0.00	0.08		
N	545	316	466	81

Notes: Sample averages across four business categories. As throughout the paper, all unrestricted variables are winsorized at the top and bottom 1 percentile.

Table A.2: Summary statistics on mechanic travel patterns

	Motor Mechanics	Food Mechanics
Travelled outside district this year (%)	20	73
Number of out-of-district trips last year	2.3	2.5
Trip distance (KM)	173	321
Customers visted	1.3	1.1
Trip revenue (USD)	51	98
Trip costs (USD)	15	32
Trip length (days)	1.4	1.9

Notes: Summary statistics of the travel behavior of mechanics in the survey. Trip costs include all costs mechanics incurred for travel, accommodation, and personal expenses and that they were reimbursed for by their customers.

Table A.3: Machine quality differences between regions

	Purchase price (USD)		Early repair frequency
	(1)	(2)	(3)
Kampala	1584.38***	1421.09***	2.07***
	(48.81)	(73.76)	(0.14)
Luwero (over Kampala)	177.67*	457.87***	0.13
	(76.06)	(91.27)	(0.21)
Napak (over Kampala)	-431.64***	-75.31	-0.90***
	(95.54)	(293.61)	(0.23)
Machine Model FE	Yes	Yes	Yes
Machine Purchase Origin FE		Yes	

Notes: Various indicators of machine quality and their differences across study regions. First two columns: dependent variable is machine purchase price when new (in USD). The average machine purchased by a Napak firm is about $400 cheaper than in Kampala, controlling for machine model (col 1). This difference becomes less pronounced once controlling for in which district the machine was purchased (column 2). Columns 3: dependent variable is the self-reported number of repairs required in the first year of operating a new machine. Machines in Napak failed roughly half as frequently, in contrast to their contemporaneous high failure rate (see Table 1).

Table A.4: Repair market costs by district and machine model

	(1) Days to repair machine	(2) Repair price (USD)
Intercept (rare model in Kampala)	3.12**	42.16***
	(1.07)	(4.04)
Machine is frequent model	1.72	-0.93
	(1.42)	(5.34)
Luwero	3.95*	4.33
	(1.92)	(7.24)
Luwero × frequent model	-3.33	2.53
	(2.36)	(8.91)
Napak	12.24**	6.26
	(4.65)	(17.54)
Napak × frequent model	-11.61*	-9.86
	(4.93)	(18.59)
N	572	572

Notes: Column (1) prints the data underlying Figure 1c. Column (2) repeats the same exercise but with an estimate of the price paid for the last repair on the left hand side. Stars indicate p-values below 0.1 (*), 0.05 (**) and 0.01 (***)

Table A.5: Gravity estimates

	(1) $\log t_{ij}$	(2) $\log \Xi_{ij}$	(3) $\log \Xi_{ij}$	(4) $\log \Theta_{ij}$	(5) $\log \Xi_{ij}$	(6) $\log \Theta_{ij}$
dur_{ij}	0.29***	-1.16**	-2.86	-0.38***	-2.02***	-0.34***
	(0.08)	(0.28)	(1.34)	(0.03)	(0.26)	(0.05)
Estimation	OLS	OLS	OLS	OLS	PPML	PPML
Destination FE	Yes	Yes	Yes	Yes	Yes	Yes
Origin FE	Yes		Yes	Yes		Yes
Implied ξ		3.98	9.80		6.91	
Implied κ_M	.29	.29	.29		.29	
Assumed θ				4.00		4.00
Implied κ_G				.1		.08
N	321	18	10	989	21	2,601

Notes: Gravity regression coefficients. Regression equations are discussed in section 5. As discussed in the main text, columns (1), (2), and (4) are the preferred estimates. Stars indicate p-values below 0.1 (*), 0.05 (**) and 0.01 (***)

Table A.6: Depreciation rates in the PWT and the Ugandan survey evidence

Asset class	**Asset subgroup**	**PWT δ_a**	**Surveyed asset**	**Ugandan $\mathbb{1}_{R,a}$**	**US $\mathbb{1}_{R,a}$**
Machinery	Computers	0.32	–	–	–
	Communication equipment	0.12	–	–	–
	Other machinery	0.13	Food proc. machines	0.22	0.07
Structures	Non-residential	0.03	–	–	–
	Residential	0.01	–	–	–
Vehicles	Transport equipment	0.19	Motorcycles	0.23	0.06
Other	Software	0.32	–	–	–
	Other intellectual property	0.15	–	–	–
	Cultivated assets	0.13	–	–	–

Notes: Depreciation rates used in the Penn World Table (PWT, Feenstra et al., 2015), as well as survey evidence constructing a comparable measure in the Ugandan sample. The PWT adopts depreciation rates δ_a for all asset classes from the US Bureau of Economic Analysis. For the purposes of constructing the capital stock, I cannot directly compare this number to my Δ measure of replacement investment, as Δ also includes things like maintenance, repair, or forgone revenue, that are not usually counted as investment. Instead, I focus on the physical "replacement" part of Δ, ie. the chance that a machine breaks enough to be fully taken out of the economy each year. Those chances are 22% and 23% each year in the Ugandan sample ("Ugandan $\mathbb{1}_{R,a}$" column), considerably higher than both the conventional depreciation rate and the US replacement analogues (last column). No comparable data is available for other asset classes.

Figure A.26: Food processing machine types



(a) Grain mill



(b) Coffee huller

C Additional details on empirical analyses.

C.1 Survey sampling procedure

The surveys were conducted in three broad geographic areas (Kampala metro area, Luwero district, and Napak district) and across four business types (commercial motorcycle riders, commercial food processing firms, motor vehicle mechanics, and food processing mechanics). The following section explains the sampling procedure and appropriate reweighting in more detail.

Geographic sampling. In Napak and Luwero district, the enumerator team was tasked to visit every parish of the district (93 in Luwero, 37 in Napak). This was achieved with the exception of one parish in Napak (Apeitilom A parish), which was not visited due to local safety concerns. Two further Napak parishes (Nariameregei and Moruongor parishes) were visited but no eligible and consenting businesses were found to include in the survey. In Kampala, 73 out of the district's 96 parishes were visited. In 22 of these, no food processing firms were sampled to avoid interfering with a concurrently running study of the grain milling industry in these parishes. The sample was instead supplemented with 36 parishes from the Kampala suburbs, which fall into neighboring Wakiso district.

In each parish, a team of two enumerators was tasked to approach every eligible food processing firm they could find. Upon consent, these were surveyed. To be eligible, a food processing firm had to be in the business of either grinding flour or hulling coffee for human consumption using machinery. In practice, over 90% of sampled food processing firms were grain mills (84% in Luwero, 95% in Kampala, and 100% in Napak). These could be businesses who either purchased unprocessed produce (grains, maize, cassava, or coffee beans), processed them on their premises and sold them on, or who waited for customers to arrive in their store and processed their produce on the spot against a fee. Figure A.26 prints photos of eligible machine types.

The team furthermore approached every food processing mechanic in a parish they could find. Since

these enterprises don't often have a storefront presence for walk-in customers but rather rely on orders from phone calls, the team asked every food processing firm for contact information of mechanics they knew. If these mechanics were based in the study area, they were contacted and, upon consent, interviewed.

There were generally too many motorcycle riders and motor vehicle mechanics in each parish to comprehensively census. Instead, the team accessed a parish map with a population-weighted centroid on their phone, ensured they went to the centroid mark, and surveyed the first two motorcycle riders and motor vehicle mechanics they could find. Only commercial motorcycle riders offering transport services were eligible, not drivers of private motorcycles. Similarly, only mechanics offering motorcycle repair services were eligible, not mechanics specialising in other motor vehicles.

Motorcycle brand oversampling. To get more power on the comparisons between different motorcycle brands, the survey team was instructed to spend an additional 10 survey days deliberately oversampling less common motorcycle makes and skipping any firms using a Bajaj brand bike. This was done after 174 motorcycles in Kampala and 143 in Luwero district had already been sampled in a representative way. I use the distribution of motorcycle makes from the representative exercise to re-weight these additionally sampled bikes appropriately. No such activity was performed in Napak.

C.2 Computing Δ

I combine survey data on the various components of replacement investments to construct an overall replacement investment rate Δ according to equation (1), restated below:

$$\Delta = \left[\underbrace{(1-\mathbb{1}_R)X}_{\substack{\text{repair + maintain} \\ \text{if not replaced}}} + \underbrace{(1-\mathbb{1}_R)W}_{\substack{\text{lost revenue} \\ \text{during repair}}} + \underbrace{(P_K - P_{K,S})\mathbb{1}_R}_{\substack{\text{discard or sell} \\ \text{and replace}}} + \underbrace{(1-\mathbb{1}_R)\delta^I}_{\substack{\text{input decay} \\ \text{if not replaced}}} \right] \times \underbrace{\frac{1}{P_K}}_{\substack{\text{normalized by} \\ \text{current value}}} \times \underbrace{\frac{1}{u}}_{\substack{\text{normalized by} \\ \text{utilization}}}$$

This section provides background on the construction and measurements of each of these components. Figure A.27 prints raw distributions of all variables going into the measure.

Replacement probability $\mathbb{1}_R$. $\mathbb{1}_R$ is one if a firm stops using a given machine over the coming year. This could be either because the firm sells the machine or discards and scraps it. The survey asks firms *"How much longer do you expect to be using this machine?"*, with the average answer being around 4 years (motorcycles) and 8 years (food processing). Answers roughly follow an exponential distribution (see panel A.27a), which I use to construct the expected probability of losing the machine in the next year. I assume machine replacement events follow a Poisson point process with parameter λ_R and that firms report their expected value of the length until the next event, or $1/\lambda_R$. I am then interested in the likelihood the next event occurs within the next 1 year. Using the CDF of the exponential distribution, I obtain

$$\mathbb{E}\,\mathbb{1}_R = 1 - \exp(-1/\text{expected years until replacement})$$

Panel A.27b prints the distribution of $\mathbb{E}\,\mathbb{1}_R$. Since respondents could only answer in entire integers, $\mathbb{E}\,\mathbb{1}_R$ by construction is bounded from above by $1 - \exp(-1) \approx 0.63$.

Repair and maintenance expenditures X. Firms report expenditures on repairs (fixing a machine after it has broken down) and maintenance (preventative activities to prolong the lifespan of a machine while it still works) in two separate ways. For repair expenditures, firms report the frequency with which a machine has broken down in the past year (δ_0, plotted in panel A.27c with a spike at $\delta_0 = 0$.), as well as the amount of money spent the last time it broke down (X_{rec}, panel A.27d).

Turning to maintenance, firms report their expenditures in two different ways. First, the survey included a module similarly to the repair costs above, asking firms how often they conduct maintenance and how much they spent on the most recent such event. In a different part of the survey, they also report how much they've spent on maintenance in total in the past month. When comparing both estimates, the former produces expenditures that are inconsistent with the latter and much too big; it is likely that firms' recollection of their last expenditure on maintenance is of a particularly thorough maintenance event, which cost more than the average such event. I hence resort to the more conservative total estimate. I take the monthly total firm expenditure on maintenance, apportion it onto the machine level using machine values as weights (in case a firm has more than one machine), and multiply by 12 to get yearly estimates. Panel A.27e prints distributions of this estimate, denoted by X_M.

Lastly, for the few firms that have a dedicated mechanic on staff, I add the mechanic's yearly wage payments, ω_M, to the estimate. Putting it all together, I arrive at

$$X = \delta_0 X_{\text{rec}} + X_M + \omega_M$$

whose distribution is printed in panel A.27f.

Lost revenue W. Similarly, I record how many days a machine laid idle and broken until it could be repaired the last time it broke, as well as whether the firm was still able to operate in the meantime (83% say they cannot). If it was able to operate, I record at what share of usual revenue. Multiplying this with the average daily revenue, I arrive at the cost of forgone revenue each time the machine breaks. Multiplying further by the annual breakdown frequency, I arrive at the total annual forgone revenue W, printed in panel A.27g. As the firm revenue distribution, W has a long right tail going up to more than \$4,000 lost by the largest firms with the longest breakage spell. For legibility, panel A.27g hence winsorizes at the 90% level (instead of the usual 1%). All calculations in the paper are performed with the full distribution, winsorized at top and bottom 1%.

Discarding and reselling replaced machines $(P_K - P_{K,S})$. After computing the replacement probability ξ, I obtain information on the likely reason for replacement, as well as whether the machine is likely to be resold and at what price. 81% state the likely reason as wanting to upgrade to a newer and better model, 13% as the machine being broken beyond repair, 6% as enterprise closure. 55% expect to sell the machine once it is discarded, on average at 70% of its current market value.

Panel A.27h prints the distribution of responses of the ratio of resale value $P_{K,S}$ to current market value P_K. The distribution has a large spike at 0 for firms that do not intend to sell on their machine once it is broken. This module was added belatedly to the survey and so only collected for roughly the last third of observations. To make progress, I use information on machine age, business sector, and the reason for discarding a machine (which was collected for all observations) and *impute* missing values of this ratio using linear projection. Panel A.27i presents the full distribution of the ratio between predicted resale price and current market value, including actual data and the 2/3 of observations using the prediction.

The contribution of machine replacement to the replacement investment rate is then, per (1):

$$\xi \frac{P_K - P_{K,S}}{P_K} = \xi \left(1 - \frac{P_{K,S}}{P_K}\right)$$

Input decay δ^I. This component captures the expenditure firms need to make throughout the year on additional intermediate inputs to achieve the same output as machines become less efficient with age. To measure this quantity I focus on fuel, likely the largest intermediate input. Firms are asked how much fuel is required to operate the machine on a typical day now (denoted f_T), and how much was required to complete a similar daily operation when the machine was first purchased (f_0).

Assuming a geometric structure of input decay in the T years since purchase, one can write

$$f_T = (\delta_{\text{dec}})^T f_0$$

where δ_{dec} is the decay rate at which the amount of required inputs grow each year. Using information on f_T, f_0, as well as the time since purchase T, I isolate the object of interest

$$\delta_{\text{dec}} = \exp\left(\log\left(\frac{f_T}{f_0}\right)/T\right)$$

whose distribution is printed in A.27j. The large spike at $\delta_{\text{dec}} = 1$ represents machines that have not gotten less efficient. Values $\delta_{\text{dec}} > 1$ represent increases in fuel requirements, the very few observations with $\delta_{\text{dec}} < 1$ (around 5% of machines) represent machines firms say have actually gotten *more* fuel efficient over the years.

To isolate the costs associated with this efficiency decay per year, I multiply yearly expenditures on intermediates C_I and materials by $(\delta_{\text{dec}} - 1)$:

$$\delta^I = (\delta_{\text{dec}} - 1) C_I$$

This procedure assumes that changes in fuel efficiency proxy well changes in overall intermediate input efficiency.

Machine market values P_K. Firms are asked to estimate the current market resale value of their machines. In a few cases (7% of observations), respondents were not able to provide an educated guess. In these cases, I use the remaining 93% of responses and predict market values non-parametrically based on machine age (using three-year age bins) and machine category (motorcycle brand or food processing machine type, such as grain mill or weighing scale). Results in this paper are robust to using only the (large) subset of firms that estimated their machine's values.

See also section C.3 right below for evidence that firms' valuations of machine resale values are reasonably accurate, at least for motorcycles where there is an active resale market.

Machine utilization u. Firms report the number of hours their machine is running and utilized on the last day the business was operational, as well as how many days in the past month, and how many months in the past year the firm was open for business. Assuming that the previous day well represents overall usage, I construct a measure of the share of the past year (consisting of 365 days and 24 hours each day) that a machine was utilized. Panel A.27k prints the distribution. On average, motorcycles

have a utilization rate of 37% (implying more than 8 hours of usage every day of the year) and food processing machines of 24% (less than 6 hours each day).

To normalize Δ by utilization, I divide the raw replacement investment rate by this measure to end up at a yearly rate at uninterrupted usage (365/24). I then scale this back down to observed levels of utilization by multiplying the resulting measure by the average usage of machines by broad machine categories (motorcycles of different brands, grain mills, weighing scales, etc.).

The final panel A.27l prints the full distribution of the replacement investment rates Δ across all firms.

C.3 Validating machine value estimates

Respondents estimate their machine's current market resale value. The exact wording of the question is *"How much would you be able to sell this motorcycle / machine for on the market today?"*. While this is a common way to probe machine values in surveys (see e.g. World Bank, 2019), concerns remain over whether respondents can accurately estimate such valuations.

To make progress, I additionally scrape data on the actual market for second-hand motorcycles in Uganda, using one of the country's largest online motorcycle marketplaces. Between February and April 2025, I scrape more than 2,500 postings of motorcycles posted all across Uganda, including machine age and make.

Figure A.6 plots the median price of a second-hand motorcycle across age-make bins on the x-axis against survey responses of motorcycle owners in the same age-make bin. I drop age-make bins with less than 5 postings on the online marketplace. Firms generally estimate prices that are reasonable and within the set of online prices, though the median guess is about 22% higher than comparable machines on the online marketplace. Respondents also correctly identify which machines are likely more or less expensive. Projecting survey estimates of each respondent i on median online prices across age-make bins b

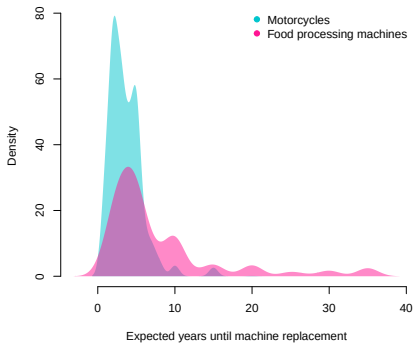
$$\text{OnlinePrice}_b = \beta_0 + \beta \text{EstimatedPrice}_{i,b} + \epsilon_{i,b}$$

reveals a projection coefficient of $\beta = 0.37$ which is highly significant.

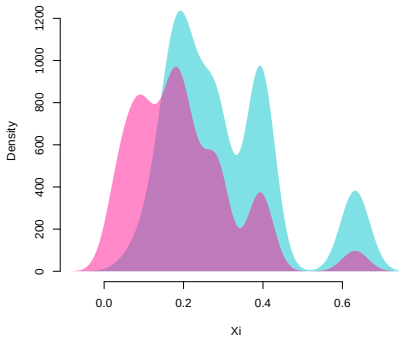
Respondents' estimates thus are at the very least reasonably accurate proxies for actual machine values. The fact that some differences between both data sources remain could be due to a variety of factors, such as (a) the online marketplace not operating in the same geographical regions as the survey activity (though respondents' estimates are not necessarily more accurate when restricting both online and survey sample to the wider Kampala region), (b) motorcycles on the resale market being on average worse quality than those in wider circulation, (c) market participants are negatively selected along dimensions such as disposable cash flow and thus willing to sell their capital at lower prices as non-participants (Lanteri, 2018), or (d) overconfidence and estimation error. Note that if the online marketplace data were correct and true machine values were indeed 22% lower, the replacement investment rate would further rise. I prefer using respondents' estimates as the more conservative approach.

Note also that no comparable data for food processing machines exists. I am thus not able to independently verify respondents' estimates in this sector.

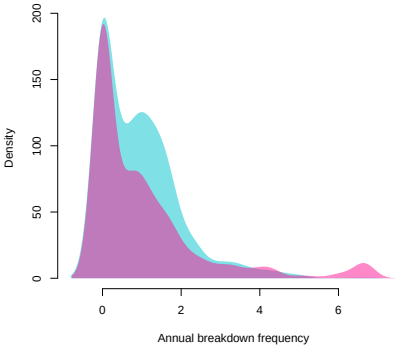
Figure A.27: Further evidence on Δ's components



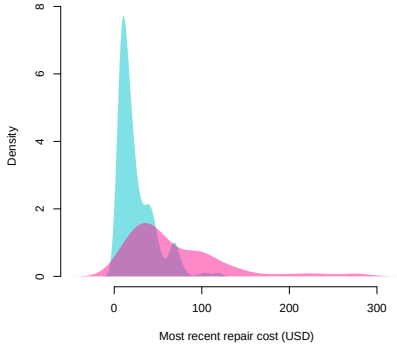
(a) Years until machine replacement



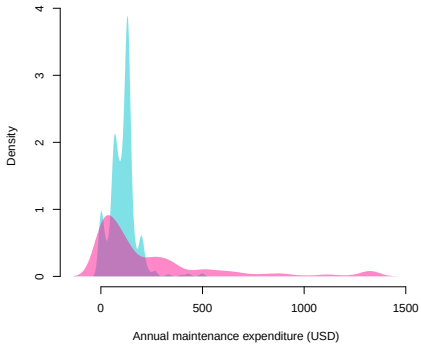
(b) Replacement probability ξ



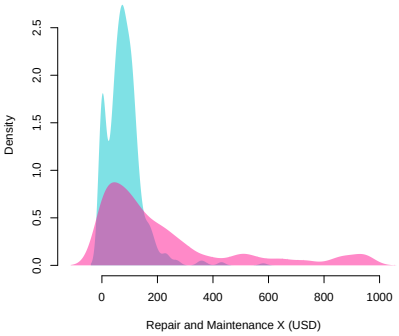
(c) Annual breakdown frequency δ_0



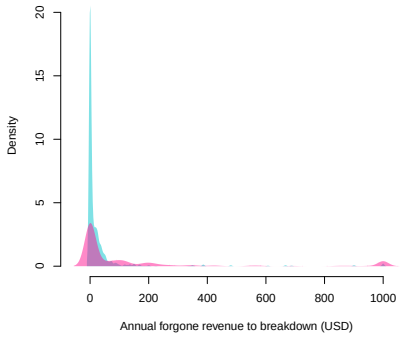
(d) Recent repair costs X_{rec} (USD)



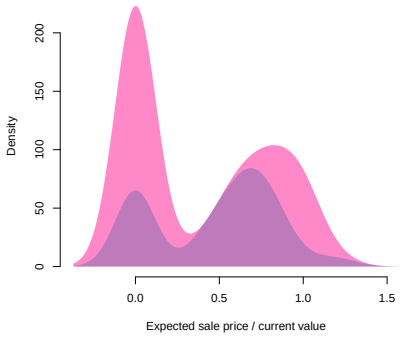
(e) Ann. maint. expenditure X_M



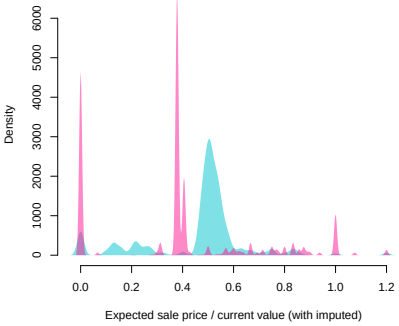
(f) Rep+maint expenditure X



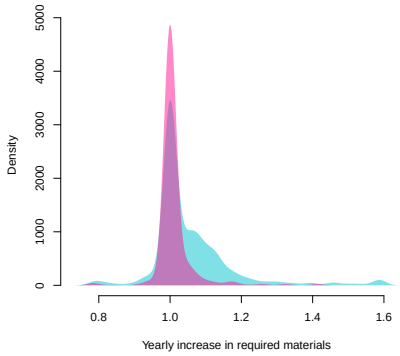
(g) Forgone rev. W (abbr. at 90%)



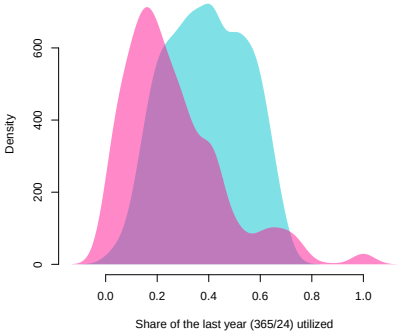
(h) Ratio resale / market value (data)



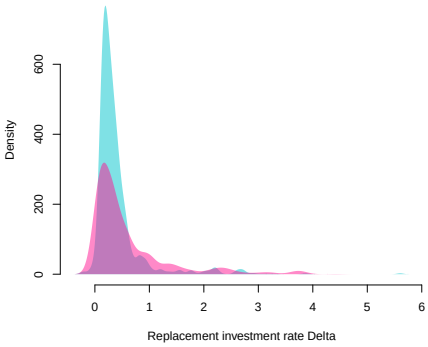
(i) Ratio incl. imputed



(j) Yearly change in fuel requ. δ_{dec}



(k) Utilization in past year (365/24)



(l) Replacement investment rate Δ

Notes: As throughout the paper, all variables are winsorized at the top and bottom 1%. The x-axis in panel A.27g is abridged at $W = 1,000$ USD for legibility, suppressing the full right tail of about 10% of observations.

C.4 Constructing replacement investment for US machines

To construct a comparable measure of replacement investment Δ for machines in the United States, I rely on a series of publicly available datasets. However, some of the components of Δ are not readily collected in existing surveys and so a full comparison would require additional data collection.

C.4.1 Farm machinery.

The American Society of Agricultural and Biological Engineers publishes engineers' estimates for the longevity, breakdown frequency, and repair costs of farm machinery in the US (ASABE Standards, 2011). Since their standards do not include an individual estimate for US grain mills, I resort to an average of the various different types of farm machinery they do cover, including sprayers, plows, harvesters, mowers, cultivators, and tractors.

Repair and maintenance. The ASABE standards report a "repair factor" for different machinery types, which is the average ratio of expected repair costs and the initial list price of a machine at various usage rates. For each machine, I consider the estimate closest to 800 hours, which is the average usage rate of the Ugandan grain mills in my sample. The average repair factor for US farm machinery used 800 hours is estimated at 16%. To account for potential differences in list price and resale price, I further divide by 0.89 which is the average such ratio in the Ugandan data (no such data exists in the US estimates), to arrive at 0.18.

Forgone revenue during idle time. The ASABE standards (Table 7.1.1) include the total expected downtime for machines in various farm activities. These are short: ranging from just 3.7 hours a year (machinery planting soybeans) to 13.6 hours a year (tillage machinery). Taking the average, I arrive at around 0.33 days of downtime. To arrive at a per-machine value, I rely on aggregate data on US farm capital-output ratios by the Iowa State University Extension School (2025). Total revenues per farm in 2023 were around $710,000 and total machinery value was $499,000. Under the conservative assumption that during the 0.33 days of downtime all revenue is lost, I arrive at a value of downtime costs per machine value of $(0.33/360) \times 710/500 = 0.0013$.

Similarly, the ASABE Standards (2011) (Table 7.1.2) publishes expected breakage rates. The published value is that the probability of *at least one* breakdown per year is around $\delta_0 = 64\%$.

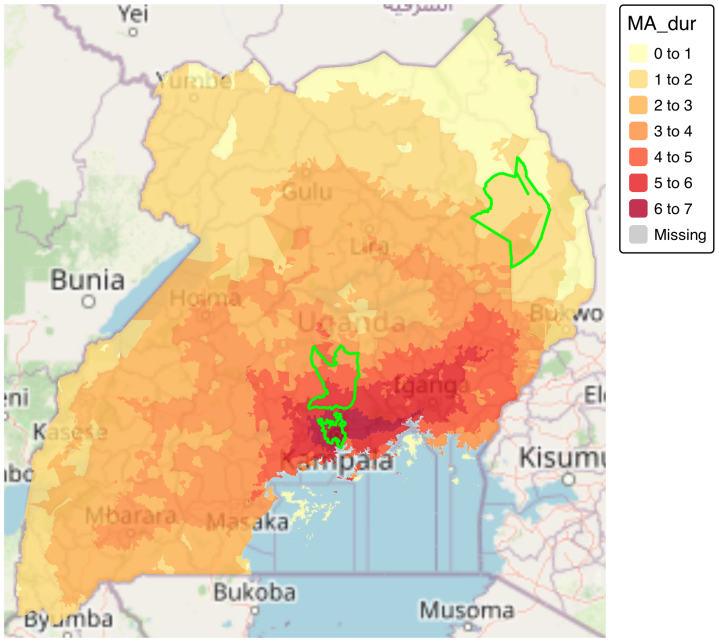
Replacement. According to the Iowa State University Extension School (2025), annual replacement rates for farm machinery in the US are around $P(replace) = 1/12$ (1/15 for tractors). The ASABE standards further publish expected salvage ratios for machines of different ages. Calibrating to the average age of food processing machines in the Ugandan sample (6 years), I arrive at an average annual replacement probability of:

$$\frac{\exp(-P(replace) \times 6)}{\sum_{n=6}^{\infty} \exp(-P(replace) \times n)} \approx 0.11$$

and an average salvage value of around 38%. The contribution to Δ is then $(1 - 0.38) \times 0.11 \approx 0.07$.

Input decay. No estimates for how much fuel efficiency declines with machinery age was available in the US.

Figure A.28: Classifying parishes by market access



Notes: Market access derived as explained in section C.5 for all 4,000+ parishes in Uganda. Study areas highlighted in green.

C.4.2 Motorcycles.

Since motorcycle transportation is not a common enterprise in the US, I focus on consumer motorcycles. Data on consumer motorcycles in the US come from the EPA and the BLS.

Repair and maintenance costs. The Consumer Expenditure Surveys in 2022 include a module on vehicle repair and maintenance expenditures. However, no direct classifier of the vehicle type is available (so expenditures for car maintenance are listed next to other vehicle types). To make progress, I use the CE vehicles module and, for each respondent household, record the vehicle type of the first, second, and third most important vehicle they own. Then I do the same for the vehicles maintenance evidence, assuming that the order in which a household answers questions about individual vehicles is the same across surveys.

With this strategy, I arrive at yearly expenditures for motorcycle repair and maintenance of $1,048. The average motorcycle in the sample is six years old and has a market value of $11,324, resulting in a contribution to Δ of around 0.09 (the market value is constructed using information on the purchase price and the BEA market price depreciation rate of 16.5%, reported in Fraumeni, 1997).

Replacement. The EPA reports that motorcycles on American roads aged 4 years (the age of machines in my survey) have an annual replacement rate of about 6%.[71] No data on replacement for resale is available.

Remaining components. Since there are basically no motorcycle transport firms in the US, forgone revenue due to motorcycles being out of operation are not immediately obtainable. Similarly, the EPA does not report any estimates for how much fuel efficiency drops per year of driving a motorcycle.

[71] See Table C-1 of the "Population and Activity of Onroad Vehicles in MOVES3" release. (link).

C.5 Classifying firms by market access

I follow Donaldson and Hornbeck (2016) and compute each location's repair market access MA_ℓ as

$$MA_\ell = \sum_{\ell'} t_{\ell\ell'}^{-\xi} \text{pop}_{\ell'}$$

at the *parish* level. There are more than 4,000 parishes in Uganda, 227 of which were visited by the survey team. I calibrate $t_{\ell,\ell'}$, the travel costs for mechanics to travel from parish ℓ' to ℓ in the same way as in the rest of the paper (using the same ξ and κ_M estimates as in Table 3). Figure A.28 prints the spatial distribution of the surrounding market access values across all parishes of Uganda. Study regions are marked in green outlays.

D Model derivations

D.1 Capitalist problem

I follow Moll (2014) and derive the solution of the capitalist's intertemporal optimization problem. To restate the model setup, capitalists hold capital K, on which they earn a return rK, but incur replacement costs ΔK each period. They can make new capital investments at exogenous price p_K. As derived in the main text, Δ is itself the outcome of an optimization problem each instant between repair, maintenance, and replacement of capital. Restating equation 2, capitalists maximize the intertemporal utility function

$$V_s = \int_s^\infty \exp(-\rho t) \log c_t dt \qquad \text{s.t.} \qquad \underbrace{p_K \dot{K}_t}_{\text{new investment}} = \underbrace{r_t K_t}_{\text{capital income}} - \underbrace{\Delta_t K_t}_{\text{replacement investment}} - \underbrace{p_{C,t} c_t}_{\text{consumption}}$$

which we can re-write as a Bellman equation:

$$\rho V(K,t) = \max_c \left\{ \log c + \frac{1}{dt} dV(K,t) \quad \text{s.t.} \quad dK = dt \left[\frac{K_t}{p_K} (r_t - \Delta_t) - \frac{p_{G,t} c}{p_K} \right] \right\} \tag{A.1}$$

As in Moll (2014), we proceed with guess-and-verify. We guess $V(K,t) = Bv(t) + B \log K$ which implies $dV(K,t) = Bdv(t) + \frac{B}{K} dK$. Plugging this into (A.1), we get

$$\begin{aligned} \rho V(K,t) &= \max_c \left\{ \log c + \frac{1}{dt} \left[Bdv(t) + \frac{B}{K} dK \right] \right\} \\ &= \max_c \left\{ \log c + \frac{Bdv(t)}{dt} + \frac{B}{K_t} \left[\frac{K_t}{p_K} (r_t - \Delta_t) - \frac{p_{G,t} c}{p_K} \right] \right\} \end{aligned}$$

forming the FOC for c yields:

$$\frac{1}{c} = \frac{B}{K_t} \frac{p_{G,t}}{p_K} \implies c^* = \frac{p_K K_t}{B p_{G,t}}$$

plugging back in, we arrive at

$$\rho V(K,t) = \log p_K + \log K_t - \log B - \log p_{G_t} + \frac{Bdv(t)}{dt} + \frac{B}{p_K} (r_t - \Delta_t) - 1$$

which implies $B \equiv 1/\rho$, and hence

$$p_{G,t}c^* = \rho p_K K \tag{A.2}$$

so capitalists always consume a constant share ρ out of their existing capital stock, which implies the LOM of capital is

$$p_K \dot{K} = K_t \left[r_t - \Delta_t - \rho p_K \right]$$

and the full value function can be written as

$$V(K,t) = \frac{1}{\rho} \left[K_t + \underbrace{\rho \log p_K + \rho \log \rho - \rho + \frac{dv(t)}{dt} - \rho \log p_{G,t} + \frac{r_t - \Delta_t}{p_K}}_{\equiv v(t)} \right]$$

implying that in steady state, $\rho p_K = r_t - \Delta_t$. Furthermore, using the constant returns to scale assumption, I can write this steady state condition as

$$\begin{aligned} rK = p_G \alpha A K^{\alpha} N_G^{1-\alpha} &= K(\Delta + \rho p_K) \\ K^* &= \left(\frac{\alpha A p_G}{\Delta + \rho p_K} \right)^{\frac{1}{1-\alpha}} N_G \end{aligned}$$

as in equation (3).

D.2 Capital-to-mechanic assignment problem

In this section, I derive the solution to the assignment problem which specifies which machine gets sent to which mechanic were it to break. I show how a certain structure of this process can microfound a continuous CES cost-minimization across mechanics.

The representative capitalist owns a mass K of machines and needs to decide to which of a mass M of mechanics to send each machine once it breaks. Index each individual machine by k and each individual mechanic by i. To guide this assignment decision, each mechanic is differentiated along three dimensions: their stated price p_i, their equilibrium processing time W_i, and a machine-by-mechanic specific cost shifter $z_k(i)$. This cost shifter is assumed to affect both the price paid p_i as well as the rental income r lost while a machine is idle in a symmetric way. This implies that a capitalist who assigned machine k to mechanic i pays $p_i \iota / z_k(i)$ as a match-specific repair price and loses $r\iota / z_k(i)$ flow rate of rental income for each instant the machine is currently sitting at mechanic i and waiting to be repaired. Recall that ι describes the severity-draw guiding how difficult this particular repair is.

Importantly, by assumption 1 capital owners make an *ex-ante* choice of where to send each machine, that is they assign each machine to a mechanic *before* it breaks and based on the long-term expected processing time. Similarly, by assumption 2 mechanics do not change their price based on how busy they currently are. This setup rules out that machines are assigned after they have already broken down and conditional on the current realization of processing times and prices. Instead, capital owners make a long-term assignment plan for each machine, which they follow in case of breakage.

Denote by $a(k,i)$ the assignment indicator, which equals one if machine k is assigned to mechanic i and

zero otherwise. Also denote by b_{ki} the probability that a unit k of capital is currently broken and waiting to be processed by mechanic i. The assignment problem of the representative capitalist is then how to choose all the $a(k,i)$ to maximize total flow capital income:

$$\max_{a(k,i)} \quad \underbrace{rK - \int_k \int_i \frac{\iota}{z_k(i)} \left(rb_{ki} + a(k,i)\delta_0 p_i\right)}_{\text{net rental income – repair costs}} \quad \text{s.t.} \quad \underbrace{\dot{b_{ki}} = \delta_0 a(k,i) - \frac{b_{ki}}{W_i}}_{\text{evolution of idle capital}} \quad \text{and} \quad \underbrace{\int_i a(k,i) = 1 \forall k}_{\text{each } k \text{ only assigned once}}$$

In words, the capitalist earns rental income r on all machines that are not currently broken, but loses $r\frac{\iota}{z_k(i)}$ rental income every instant a machine k sits at mechanic i. Similarly, the capitalist has to pay $p_i \frac{\iota}{z_k(i)}$ in monetary repair costs to mechanic i whenever the machine breaks, which happens at rate δ_0. The likelihood that machine k sits at mechanic i evolves according to the stated law of motion. Expected new failures are δ_0 if the machine is assigned to this mechanic, and expected reinstatements are $\frac{b_{bki}}{W_i}$ where $1/W_i$ is the Poisson rate at which mechanic i processes tasks entering their store after potentially waiting.

To solve this problem, we set up the Hamiltonian

$$\mathcal{H} = rK - \int_k \int_i \frac{\iota}{z_k(i)} \left(rb_{ki} + a(k,i)p_i\delta_0\right) + \eta_{ki} \int_k \int_i \left(\delta_0 a(k,i) - \frac{b_{ki}}{W_i}\right) - \int_k \varrho_k \left(\int_i a(k,i) - 1\right)$$

where η_{ki} are the coefficients on the co-state and ϱ_k are the coefficients on the second constraint. We take first order and co-state conditions:

$$\begin{aligned} \frac{\partial \mathcal{H}}{\partial a(k,i)} &= -\frac{p_i \delta_0 \iota}{z_k(i)} + \eta_{ki}\delta_0 - \varrho_k = 0 \\ \frac{\partial \mathcal{H}}{\partial b_{ki}} &= -\frac{r\iota}{z_k(i)} - \eta_{ki}\frac{1}{W_i} + \dot{\eta_{ki}} - \rho\eta_{ki} = 0 \\ \implies -\varrho_k &= \delta_0 \iota \left(\frac{p_i + rW_i}{z_k(i)}\right) \end{aligned}$$

Because of the linearity in ϱ_k, this implies that every machine is sent to the mechanic that minimizes its respective $\left(\frac{p_i + rW_i}{z_k(i)}\right)$-term:

$$a(k,i) = \begin{cases} 0, & \text{if } i \neq \arg\min_s \frac{p_s + rW_s}{z_k(s)}, \\ 1, & \text{else.} \end{cases}$$

I now make an important distributional assumption about the cost shifters $z_i(k)$, which goes back to Dupuy and Galichon (2014):

Assumption A.1 *(Repair Assignment, following Dupuy and Galichon, 2014) For each machine i, capitalists only consider a finite random subset of mechanics $i \in \mathbb{N}$ of the full mechanic mass M, whose observable attributes $x_i = (p_i + rW_i)$ and costs shifters $z_k(i)$ are points of a Poisson process of intensity $dx \times \frac{\sigma-1}{M} z^{-\sigma} dz$.*

Assumption A.1 helps making this problem more tractable. The following proof largely relies on Dupuy and Galichon (2014), who build on Dagsvik (1994), and which I adapt for the case of cost-minimization (instead of match-utility maximization). Denote $x_i \equiv p_i + rW_i$ for readability and recall that assumption 1 states that for each machine, the capitalist only considers a finite but random "plausible" subset of the infinite continuum of mechanics available. Assumption A.1 now implies that their characteristics

$(x_i, z_k(i))$ are the result of a Poisson point process with intensity $dx \times \frac{\sigma-1}{M} z^{-\sigma} dz$. What this means is that the chance that for capital unit k there exists a plausible mechanic with attributes within an infinitesimal small neighborhood dx, dz around (x, z) is equal to $\frac{\sigma-1}{M} z^{-\sigma} dx dz$, where σ is a parameter and M is the mass of mechanics in the market. The correction by the mass M of mechanics is not necessary for this proof, but chosen so that the resulting price index does not feature love-of-variety effects, which by themselves lead to agglomeration effects orthogonal to the main scale effects channel proposed by the model.

This Poisson point process also implies that the *void probability* of any set A (ie. the chance that there is no point within this set) is

$$P(\text{no point in } A) = \exp\left(-\int_A \frac{\sigma-1}{M} z^{-\sigma} dz dx\right)$$

This assumption rules out the otherwise problematic result that with infinitely many options, at least one of them should have a $z_k(i)$-term so large that the expected costs would be zero almost surely. Instead, the capitalist for each machine picks the mechanic who minimizes

$$\min_s \{\iota x_s / z_k(s)\}$$

Denote the value of this minimum by Z and note that

$$P(Z \leq c) = P(z_k(s) \leq \iota x_s / c \quad \forall s)$$

is exactly the void probability of the set $S = \{(y, z) : z_k(s) > x_s / c\}$, which collapses to

$$\begin{aligned} P(Z \leq c) &= \exp\left(-\int_S \frac{\sigma-1}{M} z^{-\sigma} dx dz\right) \\ &= \exp\left(-\int_i \int_{\iota x_s / c} \frac{\sigma-1}{M} z^{-\sigma} dx dz\right) \\ &= \exp\left(-c^{\sigma-1} \frac{\iota}{M} \int_i x_i^{1-\sigma} di\right) \end{aligned}$$

which notably is a Frechet distribution with scale $\frac{\iota}{M} \int_i x_i^{1-\sigma} di$ and shape $\sigma - 1$. Following standard Frechet algebra, the expected value of this distribution (ie. the expected minimum price paid to the chosen mechanic) is

$$\tilde{X} = \iota M^{\frac{1}{\sigma-1}} \left(\int_i x_i^{1-\sigma} di\right)^{\frac{1}{1-\sigma}} \equiv \iota X$$

which is a continuous CES price index over all mechanics with elasticity σ, as in equation (4). Similarly, the chance that $a(k, i) = 1$ is proportional to $\propto x_i^{-\sigma}$. The total market share a mechanic with prices and processing times adding up to x_i can expect (ie. its residual demand curve) is

$$\lambda_i(x_i) = \int_k a(k, i) dk = \zeta x_i^{-\sigma}$$

where $\zeta \equiv \delta_0 \iota K X^\sigma M^{-1}$. Note that in the full model, I also allow capital owners to forgo the repair market entirely in case of particularly severe ι-draws. Footnote 39 as well as the following section D.3 derive the term Ψ, which denotes the expected ι draw that actually lead to repairs. In the full model,

the demand shifter is hence $\zeta = \delta_0 \Psi K X^{\sigma} M^{-1}$.

Importantly, note that this setup ensures that the rate at which newly broken machines arrive at each mechanic λ_i is Poisson. This is because machines k are ex-ante assigned to a mechanic i according to the assignment indicators $a(k,i)$. Once these machines break (which happens at a Poisson rate), the assignment is followed, which ensures that λ_i is also Poisson. Mechanics can increase their market share by lowering prices and processing times (and hence ensuring that they are assigned more machines). Because the rate at which mechanics can solve repair tasks $n\mu$ (endogeneised just below) is also Poisson, the resulting processing rates $1/W$ are Poisson as well, justifying the law-of-motion at which capital units leave mechanics in the first place.

D.3 Optimal repair-replace decision

Once a machine breaks down, capitalists can either replace it with a new machine at price p_K (the same price as new investments), or repair it. The costs of repair depend on the severity of the breakdown. As explained in the main text, I assume each breakdown generates a *severity* ι, drawn from an exponential distribution with shape parameter γ, so that $\iota \sim \text{Exp}(\gamma)$. Optimally assigning a breakdown with severity ι to mechanics leads to a cost to the capitalist of ιX (consisting of prices and forgone rental income, derived above), determined endogenously via the repair market. This implies machines are repaired iff

$$\begin{aligned}\Pr(\text{repair}) &= \Pr(X\iota \leq p_K) \\ &= \Pr(\iota \leq p_K/X) \\ &= 1 - \exp(-\gamma p_K/X)\end{aligned}$$

so that the probability an issue is repaired goes to 1 if X becomes smaller or p_K becomes larger. High values of γ imply a lower chance of particularly severe issues (or ι-draws), which show up as a high elasticity between the two options. The estimated parameter of γ is smaller than 1, which implies capitalists sometimes replace, even though X is on average less than p_K, reflecting the real chance that some breakdowns are severe enough to not be economical to fix.

To derive the total unconditional expenditures on either of the two options, we form

$$\begin{aligned}\mathcal{R}(X) &= \mathbb{E}_\iota \min\left[p_K; X\iota\right] \\ &= p_K \times (1 - \Pr(\text{repair})) + X \int_0^{p_K/X} \iota g(\iota) d\iota \\ &= p_K \exp(-\gamma p_K/X) + X \int_0^{p_K/X} \iota \gamma \exp(-\gamma\iota) d\iota \\ &= p_K \exp(-\gamma p_K/X) + X \left[\frac{1 - (\gamma p_K/X + 1)\exp(-\gamma p_K/X)}{\gamma}\right] \\ &= \frac{X}{\gamma}\left(1 - \exp\left(\gamma p_K/X\right)\right)\end{aligned}$$

as in equation (6). This goes up in X and p_K (higher repair or replace costs mean higher costs either way) and goes down in γ (lower chance of severe issue draws).

Lastly, the expected severity draw that is actually fixed (taking into account that some particularly

severe breakages will never make it onto the mechanic market) is

$$\begin{aligned}\Psi &= \int_0^{p_K/X} \iota g(\iota) d\iota \\ &= \frac{1-(\gamma p_K/X+1)\exp(-\gamma p_K/X)}{\gamma}\end{aligned}$$

as in footnote 39.

D.4 Optimal maintenance decision

Capitalists optimally invest in maintenance m to reduce the rate at which their machines break δ_0. I assume this relationship to take the functional form

$$\delta_0(m) = \frac{1}{\nu} m^{-\nu}$$

so that more maintenance reduces the failure rate at elasticity ν. Once a machine does break, the costs to the capitalist are $\mathcal{R}$, which in itself is the outcome of the nested repair-replace problem from just above. Lastly, I assume that maintenance m can be purchased at local mechanics at linear price $c_m m$, so that the cost minimization problem can be written as

$$\begin{aligned}\min_m \Delta(m) = \delta_0(m)\mathcal{R} + c_m m &= \frac{\mathcal{R}}{\nu} m^{-\nu} + c_m m \\ \implies c_m &= (m^*)^{-\nu-1}\mathcal{R} \\ \implies \Delta(m^*) &= \frac{\mathcal{R}}{\nu}(m^*)^{-\nu} + (m^*)^{-\nu}\mathcal{R} \\ &= \mathcal{R}\frac{\nu+1}{\nu}(m^*)^{-\nu} = (\nu+1)\mathcal{R}\delta_0(m^*)\end{aligned}$$

where

$$m^* = \left(\frac{\mathcal{R}}{c_m}\right)^{\frac{1}{1+\nu}}$$

In the absence of reliable data on c_m, I instead interpret the observed $\delta_0(m^*)$ as a sufficient statistic which already solves the above problem. I estimate ν using data on the full replacement investment rate $\Delta(m^*)$ and the rate excluding maintenance expenditures $\mathcal{R}\delta_0(m^*)$:

$$\nu = \frac{\Delta(m^*)}{\mathcal{R}\delta_0(m^*)} - 1 = \frac{0.453}{0.299} - 1 = 0.51$$

D.5 Queuing behavior

The model setup is equivalent to a standard M/M/1 queue with inflow rate λ and outflow rate μn. I replicate a standard result from queuing theory that such a queue has an expected wait time

$$W = \frac{1}{\mu n - \lambda}$$

First, derive the expected number of tasks currently in the system L (these are all tasks both currently waiting to be repaired, plus the one task at the front of the line that is currently being worked on). For this, note that in steady-state, the probabilities π_j that there are currently j people in the system are

linked by the Markov chain

$$\lambda \pi_j = \mu n \pi_{j+1}$$

as otherwise the queue would explode or implode. Backwardly solving this for all $j = 0, 1, 2...$:

$$\pi_j = \left(1 - \frac{\lambda}{\mu n}\right)\left(\frac{\lambda}{\mu n}\right)^j$$

for $\mu n > \lambda$. This means the expected number of tasks currently in the system is

$$L = \sum_{j=0}^{\infty} j\pi_j = \left(1 - \frac{\lambda}{\mu n}\right)\frac{\left(\frac{\lambda}{\mu n}\right)}{\left(1 - \frac{\lambda}{\mu n}\right)^2} = \frac{\lambda}{\mu n - \lambda}$$

The expected number of outstanding orders can be linked to the expected wait time by Little's Law, which states that $W = L/\lambda$ and delivers the result.

D.6 Mechanic problem

Monopolistically competitive mechanics decide how many workers to hire n and how much to charge for their services p to maximize total profits.

$$\max_{n,p} \quad \lambda\left(p - p_s\right) - \omega n, \qquad \text{s.t.} \qquad \lambda = \zeta\left(p + \frac{r}{\mu n - \lambda}\right)^{-\sigma}$$

Note that demand inflow rate λ is a fixed-point (as more demand leads to longer queues, in turn reducing demand, and so on). We can nevertheless make progress. Take first-order conditions:

$$\begin{aligned} \frac{\partial}{\partial n}: \quad & \frac{\partial \lambda}{\partial N}(p - p_s) - \omega = 0 \\ \frac{\partial}{\partial p}: \quad & \frac{\partial \lambda}{\partial p}(p - p_s) + \lambda = 0 \\ \implies & -\frac{\partial \lambda / \partial p}{\partial \lambda / \partial n} = \frac{\lambda}{\omega} \end{aligned} \tag{A.3}$$

Now using $\lambda = \zeta\left(p + \frac{r}{\mu n - \lambda}\right)^{-\sigma}$, define:

$$\begin{aligned} A &:= \lambda\left(p + \frac{r}{\mu n - \lambda}\right)^{\sigma} - \zeta = 0 \\ \frac{\partial A}{\partial p} &= \sigma\lambda\left(p + \frac{r}{\mu n - \lambda}\right)^{\sigma - 1} \\ \frac{\partial A}{\partial \lambda} &= \left(p + \frac{r}{\mu n - \lambda}\right)^{\sigma} + \sigma\left(p + \frac{r}{\mu n - \lambda}\right)^{\sigma - 1}\frac{r}{(\mu n - \lambda)^2} \\ \frac{\partial A}{\partial n} &= \sigma\lambda\left(p + \frac{r}{\mu n - \lambda}\right)^{\sigma - 1}\frac{-r\mu}{(\mu n - \lambda)^2} \end{aligned}$$

Using the implicit function theorem:

$$\frac{\partial \lambda}{\partial N} = -\frac{\partial A/\partial n}{\partial A/\partial \lambda} = \frac{\frac{\sigma \lambda r \mu}{(\mu n - \lambda)^2}}{\frac{\sigma r}{(\mu n - \lambda)^2} + p + \frac{r}{\mu n - \lambda}}$$
$$\frac{\partial \lambda}{\partial p} = -\frac{\partial A/\partial p}{\partial A/\partial \lambda} = \frac{-\sigma \lambda}{\frac{\sigma r}{(\mu n - \lambda)^2} + p + \frac{r}{\mu n - \lambda}}$$

plugging this back into (A.3) we arrive at:

$$\frac{\lambda}{\omega} = \frac{(\mu n - \lambda)^2}{r\mu}$$
$$\implies n^* = \frac{\lambda}{\mu} + \sqrt{\frac{r\lambda}{\omega \mu}}$$

which is equation (9). Note again how this wait time *decreases* in overall demand λ, and how total costs are a *concave* function of λ. Plugging back into the mechanic problem, we get resulting wait times and marginal costs from equation (10):

$$rW = \sqrt{\frac{r\omega}{\mu \lambda}}$$
$$C(\lambda) = \lambda \left(\frac{\omega}{\mu} + p_S \right) + \sqrt{\frac{\omega \lambda r}{\mu}}$$
$$mc(\lambda) = \frac{\omega}{\mu} + p_S + 2\sqrt{\frac{\omega r}{\mu \lambda}}$$

Given the above, monopolistically competitive mechanics decide which composite cost x to charge. Note that $x = p + rW$, and hence the problem can be written as:

$$\max_x \quad \zeta x^{-\sigma} (x - rW) - C(\zeta x^{-\sigma}) = \zeta x^{1-\sigma} - \zeta x^{-\sigma} \left(\frac{\omega}{\mu} + p_S \right) - 2\sqrt{\frac{\omega \zeta x^{-\sigma} r}{\mu}} \tag{A.4}$$

The solution to this problem is a higher-order polynomial in x, which can be readily solved with available numerical methods. Section D.7 below presents an analytical solution for the case of $\sigma = 4$.

D.7 Analytic solution of the mechanic problem

For the case of $\sigma = 4$ (not far from the estimated $\sigma = 4.16$), there exists an analytical solution of the mechanic problem:

$$\max_x \quad \zeta x^{-\sigma} (x - rW) - C(\zeta x^{-\sigma}) = \zeta x^{1-\sigma} - \zeta x^{-\sigma} \left(\frac{\omega}{\mu} + p_S \right) - 2\sqrt{\frac{\omega \zeta x^{-\sigma} r}{\mu}}$$
$$\implies x^* = \mathcal{S} + \frac{2\mathcal{A}\mathcal{S}}{3\mathcal{A} + \sqrt{9\mathcal{A}^2 - 64\mathcal{S}\mathcal{C}}}$$

where

$$\mathcal{A} \equiv \zeta^{1/4}$$
$$\mathcal{C} \equiv \sqrt{\frac{r\omega}{\mu}}$$
$$\mathcal{S} \equiv p_S + \frac{\omega}{\mu}$$

First, note that in the absence of timeliness concerns (achieved through $\mu = \infty$ or $r = 0$, which in this context only refers to the cost of waiting and should not be taken literally as a zero rental rate of capital), $\mathcal{C} = 0$ which implies

$$x^* = \frac{4}{3}S$$

which is the standard CES constant-markup solution for $\sigma = 4$. Second, the same is true in the partial-equilibrium limit of infinite (and hence perfectly predictable) demand:

$$\lim_{\zeta \to \infty} = \frac{4}{3}S$$

D.8 In-house repair problem

Setup. Firms repairing their machines in-house face a similar tradeoff to the mechanic problem above: hiring a lot of repair staff assures fast repair and little downtime but comes at the cost of these employees sitting idle when not needed.[72]

In particular, I assume that capitalists in a given region rent out capital at two different rates, r_+ and r_-, where r_+ represents an all-inclusive rate with capitalists covering machine replacement investment, and r_- represents the (lower) rate for firms who are willing to perform their own in-house repair. Because capital markets are perfectly competitive, the following relationship holds:

$$r_+ = r_- + \Delta$$

so that capitalists are indifferent between both contracts. Firms are heterogeneous and produce different varieties of the final good. Consumers have CES preferences with parameter χ over varieties:

$$G = \left(\int_\varphi y(\varphi)^{\frac{\chi-1}{\chi}} d\varphi \right)^{\frac{\chi}{\chi-1}}$$

Firms draw a productivity $\varphi \sim \text{Pareto}(A, \eta)$ and produce with the production function

$$y(\varphi) = \varphi k^\alpha n^{1-\alpha}$$

Firms decide whether to rent capital using the higher rate r_+ or r_-, where the latter draws additional costs from having to cover their own replacement investment, which consists of hiring their own repair staff, paying for their own spare parts, etc. In particular, the costs associated with both arrangements

[72] As in the rest of the model, I assume repair workers cannot be used for other steps of the production process when not needed. This assumption is empirically questionable: only 25% of the firms with in-house mechanics in my survey report that these staff only conduct repair tasks, the remaining 75% also conduct production and/or organizational tasks.

are:

$$C_+(k) = r_+ k$$
$$C_-(k) = r_- k + \min_{n_R} \left\{ \delta_0 k \left(p_S + \frac{r_+}{n_R \mu - \delta_0 k} \right) + \omega n_R \right\}$$

where n_R is the number of repair workers the firm hires. By a similar argument to above, one can derive the optimal n_R and resulting costs as

$$n_R^* = \frac{\delta_0 k}{\mu} + \sqrt{\frac{\delta_0 k r_+}{\mu \omega}}$$
$$C_-^*(k) = r_- k + (1+\nu)\delta_0 k \left(p_S + \frac{\omega}{\mu} \right) + (2+\nu)\sqrt{\frac{\delta_0 k r_+ \omega}{\mu}}$$

which implies that firms should begin to rely on in-house mechanics once marginal costs under this regime are lower than relying on the repair market.

CES solution for small firms. Firms below the above cutoff are solving a standard CES monopolist problem with constant marginal costs, which we know is solved at

$$mc_+(\varphi) = \frac{\Gamma_0 r_+^\alpha \omega^{1-\alpha}}{\varphi}$$
$$p_+(\varphi) = \frac{\chi}{\chi - 1} \frac{\Gamma_0 r_+^\alpha \omega^{1-\alpha}}{\varphi}$$
$$k_+(\varphi) = \zeta \left(\frac{\chi \Gamma_0}{\chi - 1} \right)^{-\chi} r_+^{-\alpha(\chi-1)-1} \omega^{(\alpha-1)(\chi-1)} \left(\frac{1-\alpha}{\alpha} \right)^{\alpha-1} \varphi^{\chi-1}$$

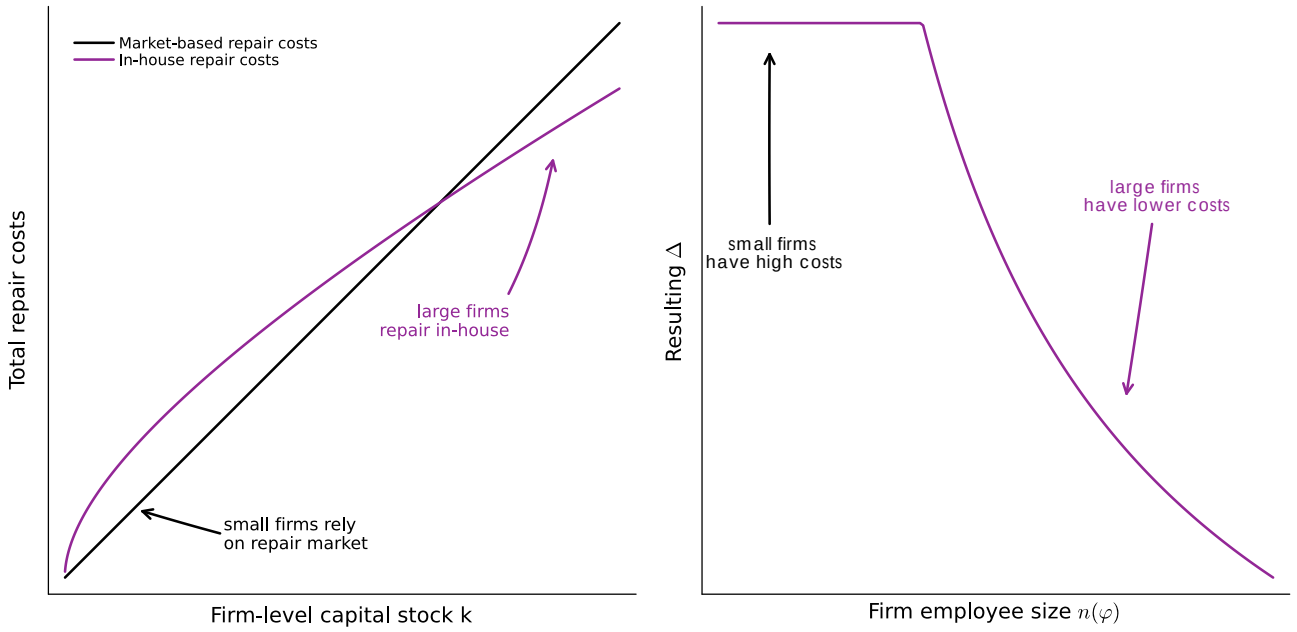
CES solution for large firms. Firms that are productive enough to find it profitable to have their own in-house repair personnel (those with $\varphi \geq \bar{\varphi}$) have a more complicated cost maximization problem, as their marginal costs decrease with scale:

$$\min \quad \omega n + C_-^*(k) \qquad \text{s.t.} \qquad y = \varphi k^\alpha n^{1-\alpha}$$

which is not generally available in closed form. Denote by $mc_-^*(\varphi)$ the resulting marginal costs from solving this above problem. Because $C_k^{*\prime} = (r_- + (1+\nu)\delta_0(p_S + \omega/\mu) + (1+0.5\nu)\sqrt{\frac{\delta_0 r_+ \omega}{\mu k}}$ decreases in k, so will mc_-^* Firms will choose in-house repair once $mc_-^* < mc_+$. Figure A.29 visualizes this intuition. Panel a shows the total costs under both regimes. The market-based repair regime has constant marginal costs to firms (in black), as every new unit of capital costs the same rental rate r_+. The in-house regime (in purple), however, again as concave total costs (or declining marginal costs), for the same reason that scale effects in the repair market exist: for low amounts of capital, there is not enough constant repair demand so that having workers on the payroll for the few times they are actually needed is not economical. However, at some point these costs flatten out and firms find it optimal to switch to the in-house regime, forgoing costly market-based repairs.[73] Panel b prints the resulting replacement investment rates Δ against firm size by number of employees n. Note the slight slight misuse of ter-

[73] Note that the costs of market-based repairs are technically taken on by capitalists, but fully passed through to firms based on the perfectly competitive nature of the rental market. High market-based repair costs are hence reflected in high marginal costs to firms based on high all-inclusive rental rates r_+.

Figure A.29: Tradeoff for firm deciding whether to hire in-house repair staff



(a) Total repair costs

(b) Resulting $\Delta \implies$ consistent with Fact #3

Notes: Repair costs and resulting replacement investment rate for firms of different sizes, once allowing for firms to endogenously decide to hire in-house repair staff. Because of returns to scale in repairs provision, only large firms will find this vertical integration attractive. Panel a reports the total repair costs from both market-based repair (as in the main model, black line) and in-house (purple line). Firms jump to repairing in-house as soon as it is advantageous to them. Panel b reports machine-level replacement investment rates Δ by firm size (which is a function of productivity φ). Larger firms resort to in-house repair, allowing them to overcome the expensive repair market, leading to a pattern qualitatively in line with stylized fact #3.

minology, as Δ for the all-inclusive package is taken on by capitalists, not firms (but perfectly passed through due to perfectly competitive capital markets). High productivity firms (which hire more n) are above the cutoff and decide to do in-house repair. They have constantly falling replacement investment rates, as larger in-house capital stocks make repairing each individual machine cheaper and faster. The simulations underlying Panel b are in line with stylized fact #3.

E Estimation details

E.1 Model inversion to back out A_ℓ

Final goods TFP parameters A_ℓ can not be readily read off observational data as this would require an external estimate of the capital stock in each region, which in practice is constructed under the assumption of a constant depreciation rate, inconsistent with the empirical findings in my paper. I hence resort to a nested routine, which for a given triple of structural parameters (σ, γ, f_E) inverts the model and backs out fundamental TFP terms A_ℓ.

In particular, I rely on the structural gravity equation (15), reprinted again below:

$$\log \Theta_{ij} = -\theta \kappa_G \text{dist}_{ij} - \underbrace{\theta \log p_j}_{f_j} + f_i$$

where again Θ_{ij} are trade import shares (the fraction of i's total consumption coming from j). The

fixed-effect on selling location j is spelled out. Using the constant returns pricing rule

$$p_\ell = \frac{\Gamma_0 \omega_\ell^{1-\alpha} r_\ell^\alpha}{A_\ell}$$

where Γ_0 again is a constant, we get that

$$\begin{aligned} f_j &= -\theta \log p_j \\ &= -\theta \log \Gamma_0 \omega_j^{1-\alpha} r_j^\alpha + \theta \log A_j \end{aligned}$$

or

$$A_j = \exp\left(f_j/\theta\right) \Gamma_0 \omega_j^{1-\alpha} r_j^\alpha \tag{A.5}$$

The fixed effects f_j are readily obtained from the gravity equation used to estimate κ_G, as explained in the main text. This suggests an estimation strategy which for a given parameter triple (σ, γ, f_E) and any starting vector of productivities $A_{\ell,0}$, I compute the resulting equilibrium quantities of ω_ℓ, r_ℓ and apply the transformation from equation (A.5) to get the next vector $A_{1,\ell}$, and so on, until convergence.

In the full estimation, I then nest this procedure inside any moment computation and fit (σ, γ, f_E), so that after this inversion, the three relevant moments are matched.