Not Worth Innovating For?
Dynamic Lock-In of Agricultural Productivity Differences*

Kazunari Tanabe†

LSE

November 12, 2026

Click HERE for the latest version.

Abstract

Why are agricultural productivity differences across countries so large? I develop a dynamic general equilibrium model in which directed innovation locks poor countries into agricultural technologies suited to rich-country environments but inappropriate for their own. Small markets can make innovation in locally adapted technologies unprofitable. These technologies then fall further behind the frontier and lose market share, eventually leading farmers to adopt inappropriate but frontier technologies. This dynamic lock-in is inefficient, because most gains from improving locally adapted technologies accrue to households and future innovators, and missing intertemporal markets prevent the innovator from appropriating them. Calibrating the model to Tanzania and the US, I find that directed innovation since 1982 contributes to their agricultural productivity difference by 42 percent in 2017. An R&D subsidy for locally adapted technologies can narrow this difference and raise welfare in Tanzania even though its households finance it.

*I am indebted to Matthias Doepke, Benjamin Moll, Rachel Ngai, and Jane Olmstead-Rumsey for their invaluable guidance and support. I also thank Kosuke Aoki, Marco Bellifemine, R. Anton Braun, Gharad Bryan, Francesco Caselli, Maarten De Ridder, Junichi Fujimoto, Maitreesh Ghatak, Stephan Hobler, Keiichiro Kobayashi, Ananya Kotia, Kiminori Matsuyama, Martí Mestieri, Dmitry Mukhin, Tommaso Porzio, Ricardo Reis, Soroush Sabet, Silvana Tenreyro, Chris Udry, and seminar participants at various venues for their helpful comments and suggestions. I am grateful to Jacob Moscona and Karthik Sastry for sharing the replication materials for Moscona and Sastry (2025). I gratefully acknowledge the generous financial support from STEG's PhD Research Grant.

†E-mail: k.tanabe2@lse.ac.uk

1 Introduction

Agricultural productivity differences across countries are enormous. In 2017, agricultural labor productivity in the richest tenth of countries was 39 times as high as in the poorest tenth, compared with a factor of 9 in non-agriculture.[1] The poorest tenth employs around 60 percent of its workers in agriculture, compared with around 2 percent in the richest tenth. Agricultural productivity is therefore central to understanding cross-country income differences (Caselli, 2005; Restuccia, Yang, & Zhu, 2008; Gollin, Lagakos, & Waugh, 2014b).

Unlike non-agricultural production, agricultural production is directly tied to local environmental conditions, such as soil, climate, and pests. As a result, the appropriateness of technology to the local environment matters far more in agriculture. Agricultural technology embodied in improved seeds, pesticides, and fertilizers works well only where the environment matches the one it is designed for. For example, pesticides developed against North American pests are less effective against pests in sub-Saharan Africa. Agricultural innovation, however, is concentrated in rich countries, and the resulting technology is designed for their environments. One possible cause of large agricultural productivity differences across countries is therefore that the technology available to poor countries is not suited to their environmental conditions. Countries whose environments differ most from those of the rich countries use the least appropriate technology, and they are also the least productive in agriculture.[2] Recent empirical work supports the view that agricultural technology developed at the frontier is inappropriate for poor countries' environments (Moscona & Sastry, 2025).

This paper makes progress in understanding the limited availability of appropriate agricultural technology for poor countries by examining how innovation incentives shape its development over time and how policy can influence this process. The inappropriateness of today's technology is itself the outcome of directed innovation in the past, and it in turn shapes the direction of future innovation. This dynamic perspective explains how inappropriate technology emerges and persists, and how its contribution to agricultural productivity differences evolves. This raises two central questions. How much does directed innovation contribute to cross-country agricultural productivity differences? Is there a role for policy intervention?

I address these questions by developing a dynamic general equilibrium model of inappropriate agricultural technology and directed innovation. Its key mechanism, dynamic lock-in

[1] The figures are calculated using the GGDC Productivity Level Database 2023 (Inklaar, Marapin, & Gräler, 2023), which covers 84 countries in 2017. Sectoral labor productivity is measured as PPP value added per person engaged in the sector. Countries are ranked by PPP aggregate value added per worker, and each reported ratio compares mean sectoral labor productivity in the top tenth with that in the bottom tenth. See also Boppart, Kiernan, Krusell, and Malmberg (2026) and Lagakos and Schoellman (2026).

[2] Among the 84 countries, the 34 whose main agricultural regions differ most from those of the US (tropical, arid, or highland, based on the dominant Köppen–Geiger climate classifications) have median agricultural labor productivity of about one thirty-fifth of the US level, compared with about one quarter for the other 50 countries. In non-agriculture, the corresponding medians are 0.28 and 0.80 of the US level.

through directed innovation, locks poor countries into inappropriate technology and amplifies agricultural productivity differences over time.

Baseline model and dynamic lock-in. To illustrate the key mechanism, I first present the baseline two-country, two-sector model. The two countries, the North and the South, differ in their agricultural environments. Agricultural technology embodied in intermediate inputs is adapted to Northern or Southern environments, with effectiveness determined by base quality and environmental mismatch. Farmers choose the input type with the higher effective quality per unit of expenditure. R&D firms in the North direct innovation toward input types whose operating profits cover the fixed cost of R&D. These profits depend on the market size for the improved inputs. Innovation in the frontier Northern type raises its base quality by a constant proportion, while innovation in the lagging Southern type draws on frontier knowledge, creating scope for larger proportional improvements and technological catch-up.

I derive three main results from this baseline model. The first is the equilibrium expression for agricultural labor productivity differences between the North and the South, which decomposes them into three sources: agricultural fundamentals; non-agricultural productivity, which affects agricultural employment through the labor-pull channel; and the ratio of the effective quality of the inputs used in the two countries. The third source is the main focus of this paper and depends on Southern farmers' technology choice. Northern farmers always use the frontier type. Southern farmers use the locally adapted type when its quality disadvantage is no greater than the frontier type's environmental mismatch penalty. Otherwise, they use the frontier type despite its mismatch. In the latter case, both countries use the frontier type, but environmental mismatch makes it less effective in the South. The contribution of this channel to agricultural productivity differences is then at its largest.

The second result is the paper's key mechanism, dynamic lock-in through directed innovation, which makes the use of inappropriate technology persistent. R&D firms improve the Southern type only if its market size, measured by the expenditure of farmers who adopt it, is large enough for operating profits to cover R&D costs. When that market size is too small, the Southern type is not improved and falls further behind as the frontier improves, which shrinks its market size further. As the quality gap to the frontier widens over time, Southern farmers eventually use the frontier type despite its environmental mismatch. Under laissez-faire, the economies remain on this path. Directed innovation thus locks Southern farmers into inappropriate technology and amplifies agricultural productivity differences over time.

The third result characterizes the dynamic amplification factor, which measures how much directed innovation amplifies agricultural productivity differences relative to a benchmark path in which both types improve every period, from the same initial state. Along the benchmark path, the Southern type improves through technological catch-up, narrowing its quality gap with the frontier, and Southern farmers continue to use locally adapted technology. Along the lock-in path, this catch-up is forgone, and the amplification factor is non-decreasing over time.

Its long-run value is greater when environmental mismatch is more severe and when greater transferability of frontier knowledge would allow the Southern type to catch up further.

Constrained efficiency. I then examine whether the decentralized equilibrium and the dynamic lock-in it generates are efficient. I consider a constrained social planner who decides, in each period and on behalf of the R&D firms, whether to improve the Southern type. The planner faces the same technological and feasibility constraints as the decentralized economy and takes the within-period equilibrium structure, including the static monopoly distortion, as given. This formulation allows me to assess whether private decisions about the direction of innovation are efficient while leaving the static monopoly distortion uncorrected.

The comparison is between the private and the social value of improving the Southern type. Under laissez-faire, the R&D firm improves it when the operating profit from selling the improved inputs covers R&D costs, so the private value is operating profit net of those costs. The social value also counts what private agents do not internalize: the continuation value of innovation, which captures the gains in later periods from the improved quality state it leaves behind. The comparison is especially clear when Southern farmers continue to use the inappropriate frontier type even after innovation, so that the improved Southern inputs attract no demand. The private value is then a loss equal to R&D costs, while the social value is the continuation value net of those costs. If the continuation value covers R&D costs, the social planner improves the Southern type, even where private firms do not. The decentralized equilibrium and the dynamic lock-in it generates are therefore constrained inefficient.

Privately optimal and constrained efficient innovation decisions differ because today's innovator cannot appropriate the future gains from the quality state its innovation creates. The intertemporal market through which these gains could be appropriated is missing. The unappropriated gains include knowledge spillovers. Improved Southern inputs raise agricultural productivity in the South and lower agricultural prices, benefiting households in both countries. Today's innovation also gives future innovators of the Southern type a better technological base. The firms that will build on the improved quality state in later periods have not yet entered, and the innovator has no claim on these future gains that it can sell. Longer patents do not eliminate this inefficiency, because a patent, however long, entitles the innovator only to the profits from its own vintage, so the innovator cannot internalize most of the continuation value, which accrues to households and to future innovators.

Quantitative model and calibration. For quantitative analysis, I extend the baseline model in three main ways. First, each input type consists of a continuum of product lines with different R&D fixed costs, so R&D firms decide not only whether to improve the Southern type but also how many of its product lines to improve. This introduces an intensive margin of directed innovation. Second, I introduce Fréchet productivity shocks across farmers and input types, so that both types are used in each country and adoption shares vary smoothly with their quality ratio.

Third, I incorporate weak intellectual property enforcement in the South and a higher fixed cost of R&D for the Southern type, both of which weaken incentives to improve it. Weak enforcement allows local imitators to capture part of innovators' profits from Southern sales, while the higher cost reflects Northern firms' limited knowledge of Southern agronomic conditions and the additional field trials needed for adaptation. The model also incorporates a wedge between agricultural and non-agricultural wages, subsistence requirements, iceberg trade costs, and land. The expression for agricultural productivity differences retains the baseline model's effective input quality channel and adds a structural change channel, which operates through agricultural employment in the South relative to the North.

I calibrate the model to Tanzania and the US, targeting moments from 2017. The three key parameters are the environmental mismatch penalty, the speed of technological catch-up, and the transferability of frontier knowledge. I calibrate the mismatch penalty using the micro estimates of Moscona and Sastry (2025), applying their estimated loss of agricultural output per unit of mismatch in crop pests and pathogens to the mismatch between Tanzania and the US. The last two are calibrated to the Green Revolution (1960-2000). The adoption shares of modern varieties in six Green-Revolution economies between 1970 and 1998, reported by Evenson and Gollin (2003), imply quality states in the model, and I choose the parameter values so that the model's law of motion fits these states most closely. The remaining inputs to quantitative analysis are the economy's quality states in 2017 and 1982, from which the policy counterfactuals and the amplification paths start. The 2017 state is measured from the varieties in farmers' fields, their ages, and the share of land under modern varieties, taken from the African Seed Access Index reports and the DIIVA variety-adoption data, and the 1982 state is inferred by running the model's law of motion backward under laissez-faire.

Quantitative results. The quantitative analysis answers the two central questions. The first answer is the dynamic amplification effect. Comparing the laissez-faire path from 1982 with the path in which both input types improve in every period, I find that directed innovation amplifies agricultural productivity differences by 42 percent in 2017 and 51 percent in the long run. The policy counterfactuals answer the second question. In each case, the economy follows the laissez-faire path from 1982 until a policy is introduced in 2017. I consider an R&D subsidy to firms improving the Southern type, designed to make its quality grow at the same rate as the frontier and thus keep the quality ratio at its 2017 value. The subsidy raises Tanzania's agricultural labor productivity relative to the US by 12 percent by 2047 and raises Southern households' welfare by 2.4 percent, even though they finance it. The welfare gain comes from cheaper food, which lowers the cost of subsistence consumption, and largely from the higher labor income of workers drawn out of agriculture and into non-agriculture by the subsidy. These income gains are valued using the measured sectoral wage gap, assuming it reflects barriers to leaving agriculture. If the gap instead reflects differences in workers' effective labor, leaving agriculture yields no income gain, and the welfare gain falls to 0.1 percent but remains positive.

Would a larger subsidy yield greater welfare gains, and could stronger intellectual property protection alone overcome dynamic lock-in? A larger R&D subsidy, designed to induce firms to improve every product line of the Southern type from 2017 onward and also financed by Southern households, raises Tanzania's agricultural labor productivity relative to the US by 52 percent by 2047 and Southern households' welfare by 6.8 percent. The welfare gain, however, depends on how the sectoral wage gap is interpreted. If leaving agriculture yields no income gain, the subsidy lowers Southern welfare by 3.2 percent. The sign depends on whether the gains from more appropriate technology cover the cost of the subsidy, and this cost is far higher when every product line is improved. Stronger enforcement of intellectual property rights in the South increases the share of the Southern type's product lines that are improved but does not alter the economy's long-run trajectory. It raises Tanzania's relative agricultural productivity in 2047 by only 0.4 percent and raises Southern households' welfare by 0.1 percent.

Related literature. First, this paper contributes to the large literature on cross-country differences in agricultural productivity (Caselli, 2005; Restuccia et al., 2008; Gollin et al., 2014b). This literature has proposed several explanations, including land and geography (Chen, 2017; Gottlieb & Grobovšek, 2019; Adamopoulos & Restuccia, 2022; Manysheva, 2024), selection of workers across sectors (Lagakos & Waugh, 2013; Herrendorf & Schoellman, 2018), transportation and trade costs (Adamopoulos, 2011; Tombe, 2015), misallocation of land across farms (Adamopoulos & Restuccia, 2014; Adamopoulos, Brandt, Leight, & Restuccia, 2022), technology adoption lags (Fujiwara & Matsuyama, 2024), capital (Chen, 2020; Caunedo & Keller, 2021), and intermediate inputs (Restuccia et al., 2008; Donovan, 2021; Boppart et al., 2026).[3,4] I contribute to this literature by modeling the R&D decisions that determine how appropriate the intermediate inputs available to poor countries are for their environments. I show that directed innovation is a dynamic source of agricultural productivity differences and that this channel is quantitatively important.

Second, the paper contributes to a broader literature on the appropriateness of technology. The idea that technologies suit the conditions for which they are developed and are less productive elsewhere goes back to Griliches (1957), Atkinson and Stiglitz (1969), and Stewart (1978), and has been developed for technologies that complement capital or skills by Basu and Weil (1998), Acemoglu and Zilibotti (2001), Caselli and Coleman (2006), and Rossi (2022).[5] I apply

[3] A related strand of the literature studies agriculture's role in economic development and structural change (e.g., Matsuyama, 1992; Gollin, Parente, & Rogerson, 2002, 2007; Alvarez-Cuadrado & Poschke, 2011; Bustos, Caprettini, & Ponticelli, 2016).

[4] A large development literature studies the barriers to the adoption of agricultural technology in poor countries (e.g., Foster & Rosenzweig, 1995; Bandiera & Rasul, 2006; Conley & Udry, 2010; Suri, 2011; Duflo, Kremer, & Robinson, 2011; Dercon & Christiaensen, 2011; Karlan, Osei, Osei-Akoto, & Udry, 2014; Emerick, de Janvry, Sadoulet, & Dar, 2016; Bold, Kaizzi, Svensson, & Yanagizawa-Drott, 2017; BenYishay & Mobarak, 2019; Beaman, BenYishay, Magruder, & Mobarak, 2021); see Foster and Rosenzweig (2010) and Suri and Udry (2022) for surveys.

[5] Relatedly, Porzio (2017) studies how the technology appropriate to a team's ability shapes the allocation of talent across countries. Acemoglu, Aghion, Barrage, and Hémous (2024) study a related mechanism, in which a

this idea to agriculture, where appropriateness depends on the local environment, in a dynamic general equilibrium model, and I quantify the role of inappropriate technology in shaping agricultural productivity differences over time.

The closest papers are Moscona and Sastry (2025) and Akerman, Moscona, Pellegrina, and Sastry (2026). Moscona and Sastry (2025) show that agricultural innovation is directed toward the ecological conditions of technology leaders and that ecological mismatch with these leaders reduces technology transfer and output. They cast the inappropriate technology hypothesis in a static model of the agricultural sector and derive their estimating equations from it. Their evidence disciplines the mismatch penalty in my calibration. My contribution has three parts. First, I model dynamic feedback between directed innovation and farmers' technology adoption, showing how it can generate dynamic lock-in and amplify agricultural productivity differences over time. Second, the model's two-sector structure captures how innovation and policy affect labor reallocation between agriculture and non-agriculture. Third, I use the two-country general equilibrium model to establish when dynamic lock-in is constrained inefficient and to evaluate how policies that change private innovators' incentives affect agricultural productivity and welfare. Akerman et al. (2026) find that Embrapa, Brazil's public agricultural research corporation, redirected innovation toward local crops and ecological conditions and substantially raised agricultural productivity. I complement their analysis of public research by evaluating policies that operate through private innovators' incentives, such as an R&D subsidy and stronger intellectual property enforcement.

Finally, the paper relates to the literature on industrial and development policy (e.g., Hausmann & Rodrik, 2003; Harrison & Rodríguez-Clare, 2010; Juhász, 2018; Itskhoki & Moll, 2019; Liu, 2019; Lashkaripour & Lugovskyy, 2023; Juhász, Lane, & Rodrik, 2024; Lane, 2025; Choi & Levchenko, 2025; Bartelme, Costinot, Donaldson, & Rodríguez-Clare, 2025; Brooks & Donovan, 2026; Kim, Lee, & Shin, 2026). The R&D subsidy I study pays R&D firms to improve technology adapted to Southern conditions, in the spirit of advance market commitments for neglected diseases (Kremer & Glennerster, 2004; Kremer, Levin, & Snyder, 2020). I contribute to this literature by deriving the rationale for such a subsidy from the dynamic inefficiency in the direction of innovation and by showing that it raises welfare in the South even when Southern households finance it.

Outline. Section 2 presents the baseline model and its key mechanism. Section 3 studies constrained efficiency. Section 4 presents the quantitative model, Section 5 its calibration, and Section 6 the quantitative results. Section 7 concludes. Proofs and derivations are in the appendices.

natural gas boom discourages clean innovation and, because the effect builds on itself over time, can prevent the energy transition.

2 Baseline Model

This section presents the baseline two-country, two-sector model of inappropriate agricultural technology and directed innovation. Agricultural technology is embodied in intermediate inputs whose effectiveness depends on their suitability for local conditions. Innovators choose which technologies to improve, while farmers choose which technologies to use. The baseline model provides a tractable setting for illustrating how their interaction generates dynamic lock-in through directed innovation. Section 4 extends it for quantitative analysis.

Section 2.1 describes the environment, Section 2.2 defines the equilibrium, and Section 2.3 characterizes agricultural productivity differences and explains the key mechanism.

2.1 Environment

Time is discrete: $t \in \{0,1,2,\ldots\}$. There are two countries, North (N) and South (S), that differ in their agricultural environments, exogenous productivities, and labor endowments. Each country has two sectors: agriculture and non-agriculture. The agricultural sector produces a good differentiated by country of origin, so the Northern and Southern agricultural goods are imperfect substitutes. The non-agricultural sector produces a homogeneous good that serves as the numéraire. Both agricultural and non-agricultural goods are freely traded between countries.

There are three types of agents. First, each sector in each country has perfectly competitive final good producers. Agricultural production uses intermediate inputs and labor, whereas non-agricultural production uses only labor. Second, there are R&D firms and intermediate input producers. R&D firms direct research toward intermediate inputs adapted to Northern or Southern environments, obtain patents for their innovations, and sell the improved inputs to farmers as monopolists. Third, each country j has a representative household consisting of a continuum of identical members of total measure L_j. Members can work as agricultural or non-agricultural workers and, in the North, as researchers for R&D firms. Labor is mobile across sectors within each country but immobile between countries.

2.1.1 Final Good Producers

Agriculture. The representative farmer in country j employs a decreasing-returns-to-scale agricultural production function:

$$Y_{j,t}^a = A_j^a \left(L_{j,t}^a\right)^\beta \left[B_{j,N} q_{N,t} x_{j,N,t} + B_{j,S} q_{S,t} x_{j,S,t}\right]^\alpha, \qquad \alpha + \beta < 1, \tag{1}$$

where $Y_{j,t}^a$ is agricultural output in country j in period t (superscripts indicate sectors: a for agriculture and n for non-agriculture), A_j^a captures exogenous agricultural productivity, reflecting factors such as inherent soil quality or management practices, $L_{j,t}^a$ denotes agricultural labor, $x_{j,k,t}$ is the quantity of intermediate inputs that are adapted to agricultural environment k and used in country j (hereafter, *type-k inputs*), and $B_{j,k} q_{k,t}$ is the effective quality of type-k

inputs in country j. The production function exhibits decreasing returns to scale in labor and intermediate inputs because a third factor, land, is in fixed supply, with a production share of $\gamma := 1 - \alpha - \beta$. The baseline model abstracts from land, which is introduced later in Section 4.[6]

The key feature of the agricultural production function (1) is that farmers use intermediate inputs x, such as pesticides, fertilizers, and improved seeds, whose effective quality depends on (i) their base quality q, determined by the innovation decisions of R&D firms, and (ii) their appropriateness for the agricultural environments in which they are used, captured by B. The North and South have distinct agricultural environments, and intermediate inputs are adapted to either Northern or Southern environments. Intermediate inputs used outside the environments for which they were designed are less effective because of environmental mismatch. For example, pests in North America are so different from those in sub-Saharan Africa that pesticides formulated for the former become less effective in the latter. Formally, the appropriateness of intermediate inputs is captured by

$$B_{j,k} := e^{-\delta \mathbb{1}\{k \neq j\}} = \begin{cases} 1 & \text{if } k = j, \\ e^{-\delta} & \text{if } k \neq j, \end{cases} \tag{2}$$

where $\delta > 0$ captures the environmental distance between the two countries, and $e^{-\delta} < 1$ is the factor by which the effective quality of an input is reduced when it is used in an environment to which it is not adapted. The effective quality of each input type in each country is given by

	type-N input	type-S input
Southern farmer	$e^{-\delta} q_{N,t}$	$q_{S,t}$
Northern farmer	$q_{N,t}$	$e^{-\delta} q_{S,t}$

In the South, the effective quality of type-N inputs is $e^{-\delta} q_{N,t}$, while that of type-S inputs is $q_{S,t}$.

The representative farmer in country j chooses agricultural labor $L^a_{j,t}$ and intermediate inputs $(x_{j,N,t}, x_{j,S,t})$ to maximize profits,

$$p^a_{j,t} Y^a_{j,t} - w^a_{j,t} L^a_{j,t} - p^x_{j,N,t} x_{j,N,t} - p^x_{j,S,t} x_{j,S,t}, \tag{3}$$

subject to the agricultural production function (1), taking as given the agricultural output price $p^a_{j,t}$, the agricultural wage $w^a_{j,t}$, and the input prices $p^x_{j,k,t}$ for $k \in \{N, S\}$. Because the two input types are perfect substitutes in effective units, the farmer uses only the input type with the highest price-adjusted effective quality at the optimum.[7] Formally, an optimal input type $\tau^*_{j,t}$ for a farmer in country j satisfies

$$\tau^*_{j,t} \in \arg\max_{k \in \{N,S\}} h_{j,k,t}, \tag{4}$$

[6] With an explicit land endowment Z_j, replace A^a_j with $A^a_j Z_j^\gamma$ in the expressions below.

[7] The two types represent alternative versions of an input that performs the same agricultural task but are adapted to different environments, such as pesticides formulated for different pest populations. Perfect substitutability in effective units provides a tractable benchmark for modeling the choice between these alternatives. It does not imply that distinct input categories, such as pesticides and fertilizers, are perfect substitutes.

where

$$h_{j,k,t} \coloneqq \frac{B_{j,k} q_{k,t}}{p^x_{j,k,t}} \tag{5}$$

measures the effective quality of type-k inputs in country j per unit of expenditure. The share of type-k inputs in country j's total intermediate input use is therefore $\omega_{j,k,t} \coloneqq \mathbb{1}\{k = \tau^*_{j,t}\}$, where $\mathbb{1}\{\cdot\}$ denotes an indicator function.[8] Thus, the baseline model features a bang-bang input choice. The optimal intermediate input demand is given by $x_{j,\tau^*_{j,t},t} = \alpha p^a_{j,t} Y^a_{j,t} / p^x_{j,\tau^*_{j,t},t}$ and $x_{j,k,t} = 0$ for $k \neq \tau^*_{j,t}$. Optimal agricultural labor demand is given by $L^a_{j,t} = \beta p^a_{j,t} Y^a_{j,t} / w^a_{j,t}$.

Non-agriculture. The representative firm in the non-agriculture sector employs a constant-returns-to-scale production function:

$$Y^n_{j,t} = A^n_j L^n_{j,t},$$

where $Y^n_{j,t}$ denotes non-agricultural output in country j, A^n_j is exogenous non-agricultural productivity, and $L^n_{j,t}$ denotes non-agricultural labor.

The firm chooses non-agricultural labor $L^n_{j,t}$ to maximize profits,

$$A^n_j L^n_{j,t} - w^n_{j,t} L^n_{j,t}, \tag{6}$$

taking the non-agricultural wage $w^n_{j,t}$ as given. The optimality condition is $w^n_{j,t} = A^n_j$.

2.1.2 R&D Firms and Intermediate Input Producers

The base quality $q_{k,t}$ of type-k inputs improves through innovation, with R&D firms directing research toward type-N or type-S inputs. Research on the frontier technology raises its base quality by a factor of $1 + \xi$ with certainty. Research on the lagging technology draws on frontier knowledge, raising its base quality by a larger factor with certainty and promoting technological catch-up. Throughout the analysis, type-N inputs represent the frontier technology, while type-S inputs represent the lagging technology. An innovator who conducts research on type-k inputs in period t receives a one-period patent and sells the improved inputs to farmers.

Timing. There is a pool of potential R&D firms located in the North.[9] R&D decisions and input production in period t proceed as follows.

Stage 1: Entry and R&D decisions. At the beginning of period t, potential R&D firms decide whether to undertake research and, if so, which input type to improve. Entry is sequential and observable, and all research builds on the best vintages available at the

[8] When $h_{j,N,t} = h_{j,S,t}$, I assume that the farmer chooses the input type adapted to the local environment, so that $\tau^*_{j,t} = j$.

[9] Appendix A.6 allows R&D in both countries and shows that it occurs only in the North in equilibrium.

end of period $t-1$. The first firm to enter and improve type-k inputs receives a one-period patent granting it the exclusive right to produce and sell the improved inputs in Stage 2.

Stage 2: Production. The patent holder becomes an intermediate input producer and sells the improved inputs to farmers during period t. If no innovation occurs for type-k inputs at the beginning of period t, a competitive fringe of input producers supplies the previous best vintage at marginal cost.

Stage 3: Patent expiry. At the end of period t, all patents granted in that period expire, and the corresponding blueprints become publicly available. From period $t+1$ onward, the competitive fringe can produce these vintages at marginal cost.

For each input type k, patent exclusion ensures that at most one R&D firm enters and undertakes research in each period. Because all firms build on the same previous-period vintages, a subsequent entrant researching the same input type would develop identical inputs but would be excluded from selling them during period t. It would therefore be unable to recover its fixed cost of R&D, deterring further entry into research on that input type.[10]

The R&D firm's optimization problem is solved by backward induction.

Stage 2: Production. Conditional on innovation in type-k inputs, the patent-holding R&D firm produces the improved type-k inputs and sells them to farmers. Producing one unit of either input type, regardless of vintage, requires ψ units of the non-agricultural good. Without loss of generality, I normalize $\psi = \alpha^2$. The R&D firm holding the patent for type-k inputs chooses prices $p^x_{j,k,t}$ for $j \in \{N,S\}$ to maximize operating profits,

$$\sum_{j\in\{N,S\}} \left(p^x_{j,k,t} - \psi\right) x_{j,k,t},$$

subject to farmers' demand for type-k inputs in each country, $x_{j,k,t} = x_{j,k,t}(p^x_{j,k,t})$. Because the competitive fringe can supply type-k inputs of base quality $q_{k,t-1}$ at marginal cost ψ, the producer of the improved inputs may need to set a limit price to prevent farmers from switching to the previous vintage. The optimal price set by the patent-holding R&D firm is:

$$p^x_{j,k,t} = \min\{\alpha,\ G_{k,t}\,\psi\}, \tag{7}$$

where $G_{k,t} \coloneqq q_{k,t}/q_{k,t-1}$ measures the size of the quality improvement from innovation in type-k inputs. The patent-holding R&D firm charges the unconstrained monopoly price α when

[10] Alternatively, suppose that independent innovators of the same vintage are permitted to produce and sell it, while non-innovators cannot copy it during period t. Under sequential, observable entry, a second innovator of the same input type would supply identical inputs at the same marginal cost. Bertrand competition would then yield zero operating profits, leaving the entrant unable to recover its fixed cost of R&D. Anticipated competition therefore deters duplicate research, ensuring that at most one firm innovates each input type in each period.

$G_{k,t} \geq 1/\alpha$. I refer to this case as drastic innovation. If $G_{k,t} < 1/\alpha$, the competitive fringe caps the firm's price at the limit price $G_{k,t}\psi$. This case corresponds to non-drastic innovation. Since the optimal price $p^x_{j,k,t}$ does not depend on j, I henceforth suppress this index and write $p^x_{k,t}$. The profit function of the type-k input monopolist is given by

$$\pi_{k,t} := \nu(G_{k,t})\left[\omega_{N,k,t} \cdot \alpha p^a_{N,t} Y^a_{N,t} + \omega_{S,k,t} \cdot \alpha p^a_{S,t} Y^a_{S,t}\right], \tag{8}$$

where

$$\nu(G) := \frac{p^x - \psi}{p^x} = \min\left\{1-\alpha,\, 1-\frac{1}{G}\right\} \tag{9}$$

is the monopolist's margin rate, defined as its operating profit per unit of revenue from intermediate input sales in the countries it serves. For convenience, I write $\nu_{k,t} := \nu(G_{k,t})$ for $k \in \{N,S\}$.

Stage 1: Entry and R&D decisions. The R&D block specifies quality improvements and fixed costs of R&D. Innovation raises an input type's base quality with certainty, and the size of the quality gain depends on whether research targets the frontier technology or the lagging technology. The frontier technology evolves as in a standard quality-ladder model: innovation in period t raises its base quality to $q_{N,t} = (1+\xi)q_{N,t-1}$, where $\xi > 0$; otherwise, $q_{N,t} = q_{N,t-1}$. Innovation in the lagging technology also incorporates catch-up:

$$q_{S,t} = \begin{cases} (1+\xi)\left[a\chi q_{N,t-1} + (1-a)\, q_{S,t-1}\right] & \text{if type-}S\text{ inputs are improved in period } t, \\ q_{S,t-1} & \text{otherwise,} \end{cases} \tag{10}$$

where the adaptation rate $a \in (0,1]$ is the weight on transferable frontier knowledge, and transferability $\chi \in (0,1]$ is the fraction of frontier quality that can be transferred to Southern environments. An R&D firm improving the lagging technology combines a fraction a of the transferable knowledge from the Northern frontier with a fraction $1-a$ of the knowledge embodied in the existing Southern vintage. It then applies the quality step $1+\xi$ to this knowledge base. This mechanism is analogous to the advantage of backwardness emphasized by Gerschenkron (1962): the further a technology lags behind the frontier, the greater its potential for technological catch-up through knowledge spillovers.

Innovating either input type requires $1/\eta$ researchers, where $\eta > 0$ denotes research productivity. The R&D fixed cost per innovation is therefore

$$c := \frac{w^r_{N,t}}{\eta}, \tag{11}$$

where $w^r_{N,t}$ denotes the wage paid to researchers in the North.

An R&D firm enters and innovates type-k inputs if its anticipated operating profit covers the fixed cost of R&D:

$$\Pi_{k,t} := \pi_{k,t} - c \geq 0, \tag{12}$$

where $\Pi_{k,t}$ is the firm's net profit from innovating and supplying type-k inputs. The profitability of each input type therefore determines the direction of innovation.

2.1.3 Households and Market Clearing

Households. The representative household in country j has preferences over the consumption aggregate given by

$$U_j = \sum_{t=0}^{\infty} \frac{1}{(1+\rho)^t} u(C_{j,t}), \qquad u(C) = \frac{C^{1-\sigma}-1}{1-\sigma}, \tag{13}$$

where $\rho > 0$ is the discount rate and the consumption aggregate $C_{j,t}$ is given by

$$C_{j,t} = \left(C^a_{j,t}\right)^{\phi} \left(C^n_{j,t}\right)^{1-\phi}, \qquad C^a_{j,t} = \left[\mu^{\frac{1}{\varepsilon}} \left(c^a_{j,N,t}\right)^{\frac{\varepsilon-1}{\varepsilon}} + (1-\mu)^{\frac{1}{\varepsilon}} \left(c^a_{j,S,t}\right)^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}}. \tag{14}$$

The consumption aggregate combines agricultural and non-agricultural consumption. Agricultural consumption is a CES aggregate of varieties differentiated by country of origin, where $c^a_{j,k,t}$ is country j's consumption of the agricultural variety produced in country k, and μ is the preference weight on the Northern variety. Throughout the analysis, I assume equal weights on the Northern and Southern varieties in both countries. The parameter ε is the elasticity of substitution between agricultural varieties. I assume $\varepsilon > 1$, so the varieties are gross substitutes.

The household's period budget constraint is given by

$$p^a_{N,t} c^a_{j,N,t} + p^a_{S,t} c^a_{j,S,t} + C^n_{j,t} = E_{j,t}, \tag{15}$$

where $E_{j,t}$ denotes household income in country j and is given by

$$\begin{aligned} E_{N,t} &= w^a_{N,t} L^a_{N,t} + w^n_{N,t} L^n_{N,t} + w^r_{N,t} L^r_{N,t} + \text{Profit}^a_{N,t} + \sum_{k \in \{N,S\}} \Pi_{k,t}, \\ E_{S,t} &= w^a_{S,t} L^a_{S,t} + w^n_{S,t} L^n_{S,t} + \text{Profit}^a_{S,t}, \end{aligned} \tag{16}$$

where $\text{Profit}^a_{j,t}$ denotes the profits of agricultural final good producers in country j. Households receive labor income and profits from the firms they own. They supply labor to agricultural producers, non-agricultural firms, and R&D firms. Households own agricultural final good producers and R&D firms located in their country. Since R&D takes place only in the North, Southern households receive neither labor income nor profits from R&D firms. There are no assets that allow households to transfer income across periods.

Market clearing. Market clearing for agricultural goods requires

$$c^a_{N,j,t} + c^a_{S,j,t} = Y^a_{j,t}, \qquad \forall j \in \{N,S\},$$

while market clearing for the non-agricultural good requires

$$C^n_{N,t} + C^n_{S,t} + \sum_{k \in \{N,S\}} \psi(x_{N,k,t} + x_{S,k,t}) = Y^n_{N,t} + Y^n_{S,t}.$$

The labor market clearing conditions are:

$$\begin{aligned} L^a_{N,t} + L^n_{N,t} + L^r_{N,t} &= L_N, \\ L^a_{S,t} + L^n_{S,t} &= L_S. \end{aligned}$$

2.2 Equilibrium

The definition of equilibrium is as follows.

Definition 1 (Equilibrium)**.** Given $(q_{N,-1}, q_{S,-1})$ with $\lambda_{-1} \in (0, \chi]$, an equilibrium consists of a sequence of prices $\left\{ \left(p^a_{j,t}, w^a_{j,t}, w^n_{j,t} \right)_{j \in \{N,S\}}, w^r_{N,t}, \left(p^x_{k,t} \right)_{k \in \{N,S\}} \right\}_{t=0}^{\infty}$ and corresponding quantities such that, in every period t, (i) farmers maximize (3) subject to (1), taking prices as given; (ii) non-agricultural final good producers maximize (6), taking prices as given; (iii) the type-k patent holder maximizes its operating profit subject to farmers' demand, and input types that are not improved are supplied by the competitive fringe at marginal cost; (iv) an R&D firm improves type-k inputs if and only if (12) holds; (v) households maximize (13) subject to (15) and (16); (vi) base qualities evolve according to the quality ladder for type-N inputs and the catch-up law (10) for type-S inputs; (vii) wages are equalized across sectors, so that $w^a_{j,t} = w^n_{j,t}$ for $j \in \{N, S\}$ and $w^r_{N,t} = w^n_{N,t}$; and (viii) the markets for the two agricultural goods, the non-agricultural good, and labor clear.

Recursive structure of the equilibrium. The endogenous state variable is the base quality ratio of type-S to type-N inputs, $\lambda_t := q_{S,t}/q_{N,t}$. The equilibrium is recursive in the inherited state λ_{t-1}. Given λ_{t-1}, the period-t equilibrium is solved, yielding the updated ratio λ_t, which then serves as the inherited state for period $t+1$. R&D firms' innovation decisions govern the law of motion for λ_t. The main focus of this paper is on whether agricultural technologies adapted to poor-country environments improve through innovation under laissez-faire, and on how this shapes the technology gap and resulting cross-country differences in agricultural productivity. To focus on this margin, I assume that the parameters are such that $\Pi_{N,t} \geq 0$ for all t in equilibrium, so that the frontier technology improves in every period. Under this assumption, the law of motion for λ_t depends on whether R&D firms innovate in type-S inputs:

$$\lambda_t = \begin{cases} T(\lambda_{t-1}) := (1-a)\lambda_{t-1} + a\chi & \text{if type-}S\text{ innovation occurs in period } t, \\ \lambda_{t-1}/(1+\xi) & \text{otherwise.} \end{cases} \tag{17}$$

The law of motion (17) implies that, with type-S innovation in every period, the quality ratio λ_t rises toward the ceiling χ set by the transferability of frontier knowledge, narrowing the quality gap. In the absence of type-S innovation, λ_t declines toward zero, widening the quality gap.

Drastic and non-drastic innovation. The patent-holding R&D firm's optimal price and the corresponding margin rate depend on the size of the quality improvement resulting from innovation. Innovation in type-N inputs has a constant gross step size, $G_{N,t} = 1 + \xi$. For the empirically relevant parameter values presented in Section 5.1.1, this step size satisfies $1 + \xi < 1/\alpha$, so type-N innovation is non-drastic and the patent holder charges the limit price. The corresponding margin rate in (9) is therefore constant at $\nu_{N,t} = \xi/(1+\xi)$. In contrast, type-S innovation

incorporates technological catch-up, so its gross step size increases as the inherited quality gap widens. When λ_{t-1} is sufficiently small, type-S innovation is drastic and the patent holder charges the unconstrained monopoly price. Otherwise, it is non-drastic and the patent holder charges the limit price. The following lemma formalizes this distinction for type-S innovation.

Lemma 1. *Suppose that $1+\xi < 1/\alpha$. If type-S innovation occurs in period t with inherited quality ratio $\lambda_{t-1} \leq \chi$, its gross step size is given by*

$$G_{S,t} = G_S(\lambda_{t-1}) := (1+\xi)\frac{T(\lambda_{t-1})}{\lambda_{t-1}} = (1+\xi)\Big[(1-a) + \frac{a\chi}{\lambda_{t-1}}\Big], \tag{18}$$

where G_S is strictly decreasing in λ_{t-1}, with $G_S(\chi) = 1+\xi$. Type-S innovation is drastic if and only if $\lambda_{t-1} \leq \lambda^D$, where

$$\lambda^D := \frac{\alpha(1+\xi)\,a\chi}{1-\alpha(1+\xi)(1-a)}. \tag{19}$$

The type-S patent holder's margin rate is given by

$$\nu_{S,t} = \nu(G_S(\lambda_{t-1})) = \begin{cases} 1-\alpha & \text{if} \quad \lambda_{t-1} \leq \lambda^D, \\ 1 - \dfrac{\lambda_{t-1}}{(1+\xi)\,T(\lambda_{t-1})} & \text{if} \quad \lambda^D < \lambda_{t-1} \leq \chi. \end{cases} \tag{20}$$

This margin rate is constant at its maximum on $(0, \lambda^D]$ and strictly decreasing in λ_{t-1} on $[\lambda^D, \chi]$. At $\lambda_{t-1} = \chi$, it equals the type-N margin rate $\nu_{N,t}$.

Proof. See Appendix A.2. □

Thus, a lower inherited quality ratio λ_{t-1} allows the type-S patent holder to raise its price and margin rate, up to their unconstrained monopoly levels. However, the associated quality disadvantage can also reduce farmers' demand for type-S inputs, potentially to zero in the baseline model. Consequently, even when $\lambda_{t-1} \leq \lambda^{\mathrm{D}}$ and innovation would be drastic, operating profits may be insufficient to cover the fixed cost of R&D, so type-S innovation need not occur.

Allocation-relevant quality ratio. The quality ratio relevant for period-t allocations can exhibit a one-period lag, depending on whether drastic type-S innovation occurs. Define

$$\widehat{\lambda}_t := \frac{q_{S,t}/p^x_{S,t}}{q_{N,t}/p^x_{N,t}}$$

as the price-adjusted quality ratio. Farmers' optimal input type choice (4) and the resulting adoption shares depend on this ratio, rather than directly on the base quality ratio λ_t. This ratio also enters the equilibrium expressions for South-to-North ratios, including the agricultural labor productivity ratio introduced in Section 2.3.1. Thus, $\widehat{\lambda}_t$ serves as the allocation-relevant quality ratio in period t. If both input types were sold at the same price, prices would cancel

from this ratio, yielding $\widehat{\lambda}_t = \lambda_t$. Under equilibrium pricing, however, prices generally differ, and

$$\widehat{\lambda}_t = \begin{cases} \lambda_{t-1} & \text{if drastic type-}S\text{ innovation does not occur in period } t, \\ \alpha(1+\xi)\, T(\lambda_{t-1}) & \text{if drastic type-}S\text{ innovation occurs in period } t. \end{cases} \tag{21}$$

Unless type-S innovation is drastic and the patent holder charges the unconstrained monopoly price, the allocation-relevant quality ratio equals the inherited base quality ratio, $\widehat{\lambda}_t = \lambda_{t-1}$. In this sense, the quality ratio relevant for current allocations exhibits a one-period lag.

2.3 Equilibrium Characterization and Key Mechanism

This section presents the equilibrium expression for agricultural productivity differences and explains the model's key mechanism: dynamic lock-in through directed innovation.

2.3.1 Agricultural Productivity Differences

The following proposition characterizes agricultural labor productivity differences between the South and the North in equilibrium.

Proposition 1. *In equilibrium, agricultural labor productivity differences between the South and the North in period t are given by*

$$\frac{Y^a_{S,t}/L^a_{S,t}}{Y^a_{N,t}/L^a_{N,t}} = \left[\left(\frac{A^a_S}{A^a_N}\right) \cdot \left(\frac{A^n_S}{A^n_N}\right)^{\alpha+\gamma\varepsilon} \cdot \mathcal{Q}\left(\widehat{\lambda}_t\right)^{\alpha} \right]^{1/\Gamma}, \tag{22}$$

where $\mathcal{Q}(\widehat{\lambda}) := \max\left\{\widehat{\lambda}, e^{-\delta}\right\}$ *denotes the ratio of the effective intermediate input quality between the South and the North, and* $\Gamma := 1 + \gamma(\varepsilon - 1)$.

Proof. See Appendix A.5. □

Equation (22) decomposes agricultural productivity differences into three sources. The first is agricultural fundamentals: exogenous differences in agricultural productivity. A country with stronger agricultural fundamentals has higher agricultural labor productivity. The second is non-agricultural productivity, which affects agricultural employment through wage equalization across sectors. Specifically, wage equalization implies that $L^a_{j,t} = \beta p^a_{j,t} Y^a_{j,t} / A^n_j$, so higher non-agricultural productivity reduces agricultural employment for a given level of agricultural revenue. By drawing labor out of agriculture, higher non-agricultural productivity reduces congestion in agriculture and raises agricultural labor productivity. This is the "labor pull" channel, whose importance is supported by several recent papers (e.g., Boppart et al., 2026; Caunedo, Felix, & Manysheva, 2026). The exponent $\alpha + \gamma\varepsilon$ captures the direct effect of the wage on the marginal product, as well as its indirect effect through the variety price.

The third is the main object of this paper: the ratio $\mathcal{Q}(\widehat{\lambda}_t)$ of effective intermediate input quality between the South and the North, raised to the intermediate input share α. This term reflects differences in the effective quality of the inputs used by farmers in the two countries in period t. In equilibrium, Northern farmers always use type-N inputs. In contrast, Southern farmers' choice between type-N and type-S inputs depends on the allocation-relevant quality gap $\widehat{\lambda}_t$. If $\widehat{\lambda}_t < e^{-\delta}$, or equivalently if

$$\frac{q_{S,t}}{p^x_{S,t}} < \frac{e^{-\delta} q_{N,t}}{p^x_{N,t}},$$

then type-N inputs have higher price-adjusted effective quality in the South, and Southern farmers therefore use them. Although environmental mismatch reduces the effectiveness of type-N inputs in Southern agricultural conditions, locally adapted type-S inputs are of too low a quality to be preferable. Southern farmers therefore use inappropriate agricultural technology. In this case, farmers in both countries use type-N inputs, so the common factor $q_{N,t}/p^x_{N,t}$ cancels from the ratio, while the environmental mismatch penalty $e^{-\delta}$ remains in the numerator, yielding $\mathcal{Q}(\widehat{\lambda}_t) = e^{-\delta}$. If $\widehat{\lambda}_t \geq e^{-\delta}$, Southern farmers use locally adapted type-S inputs. In this case, $\mathcal{Q}(\widehat{\lambda}_t)$ is the ratio of the price-adjusted quality of type-S inputs used in the South to that of type-N inputs used in the North, yielding $\mathcal{Q}(\widehat{\lambda}_t) = \widehat{\lambda}_t$.

2.3.2 Dynamic Lock-in through Directed Innovation

The previous section characterized agricultural productivity differences in period t, given the inherited quality state λ_{t-1}. The interaction between R&D firms' directed innovation and farmers' adoption decisions governs both the period-t value of $\mathcal{Q}(\widehat{\lambda}_t)$ and its evolution. This section examines how this interaction locks Southern farmers into inappropriate technology and amplifies agricultural productivity differences over time.

Dynamic lock-in. R&D firms direct innovation toward input types that are profitable to improve. The profitability of type-k innovation depends on its market size: how much farmers would spend on improved inputs. For type-S inputs, this demand depends on their quality relative to frontier type-N inputs and on agricultural fundamentals in the countries where they would be used. When type-S inputs lag far behind the frontier, demand for improved inputs can be too weak for innovation to cover R&D fixed costs, despite the potential for technological catch-up. Without type-S innovation, these inputs fall further behind as frontier type-N inputs improve. This can create a feedback loop in which weak demand discourages type-S innovation, while the absence of innovation widens the quality gap and keeps demand weak.

In the model, this feedback loop is captured by the equilibrium law of motion for the quality ratio λ_t. Given the inherited quality state λ_{t-1}, R&D firms decide whether to improve type-S

inputs based on profitability. Operating and net profits from type-S innovation are

$$\pi_{S,t}(\lambda_{t-1}) = \underbrace{\nu_{S,t}}_{\text{margin rate}} \underbrace{\left[\omega_{N,S,t} \cdot \alpha p^a_{N,t} Y^a_{N,t} + \omega_{S,S,t} \cdot \alpha p^a_{S,t} Y^a_{S,t}\right]}_{\text{market size}} \quad \text{and} \quad \Pi_{S,t}(\lambda_{t-1}) = \pi_{S,t}(\lambda_{t-1}) - c.$$

Market size depends on two components in each country: the adoption share $\omega_{j,S,t}$ and total expenditure on intermediate inputs, $\alpha p^a_{j,t} Y^a_{j,t}$. The adoption share determines how much of this expenditure is allocated to type-S inputs, while total expenditure determines the size of the market available to the innovator. Adoption decisions depend on the price-adjusted quality ratio $\widehat{\lambda}_t$, which is a function of the inherited state λ_{t-1} (see equation (21)). Total expenditure depends on agricultural fundamentals and the quality of the inputs used. Operating profits also depend on the margin rate $\nu_{S,t}$. A lower inherited quality ratio λ_{t-1} provides greater scope for technological catch-up and can therefore support a higher margin rate.

In equilibrium, Northern farmers always use type-N inputs, so $\omega_{N,S,t} = 0$ and the market for type-S inputs is confined to the South. If $\widehat{\lambda}_t < e^{-\delta}$ even after the quality improvement from type-S innovation, Southern farmers also choose type-N inputs, so $\omega_{S,S,t} = 0$. In this case, there is no demand for type-S inputs and $\Pi_{S,t}(\lambda_{t-1}) = -c$. Technological catch-up may allow the innovator to charge a higher price and earn a higher margin rate, but a higher margin rate cannot generate operating profits without demand. If instead $\widehat{\lambda}_t \geq e^{-\delta}$, Southern farmers adopt type-S inputs, so $\omega_{S,S,t} = 1$. The market size then equals Southern farmers' total expenditure on intermediate inputs, $\alpha p^a_{S,t} Y^a_{S,t}$. Even with adoption, this market may be too small to support innovation. Specifically, if $\nu_{S,t} \alpha p^a_{S,t} Y^a_{S,t} < c$, operating profits fall short of the fixed cost of R&D, so $\Pi_{S,t}(\lambda_{t-1}) < 0$ and type-S innovation does not occur.

The innovation decision determines which branch of the law of motion (17) applies. When $\Pi_{S,t}(\lambda_{t-1}) < 0$, type-$S$ inputs are not improved, while type-N inputs improve. The quality ratio therefore falls to $\lambda_t = \lambda_{t-1}/(1+\xi) < \lambda_{t-1}$. This lower ratio becomes the state for period $t+1$.

The lower inherited quality ratio feeds back into Southern farmers' adoption decisions in period $t+1$, so the market for type-S inputs can shrink or remain absent. If $\widehat{\lambda}_{t+1} < e^{-\delta}$ even after innovation, demand is zero and $\Pi_{S,t+1}(\lambda_t) = -c$. Even with adoption, the market may still be too small to support innovation. If $\Pi_{S,t+1}(\lambda_t) < 0$, type-$S$ innovation again does not occur, and the quality ratio falls further to $\lambda_{t+1} = \lambda_t/(1+\xi)$. Once the economies enter this loop, they remain on a path without type-S innovation under laissez-faire. Directed innovation thus dynamically locks Southern farmers into inappropriate type-N inputs, and the resulting environmental mismatch contributes to long-run differences in agricultural productivity. This feedback loop is the paper's main mechanism: dynamic lock-in through directed innovation.

Innovation with a sufficiently large type-S market size. To highlight the role of market size in dynamic lock-in, consider a special case with symmetric economic fundamentals. Specifically, suppose that the two countries have identical economic fundamentals ($A^a_S = A^a_N$, $A^n_S = A^n_N$, and $L_S = L_N$) and that the inherited quality ratio λ_{t-1} is not too low, so that $\widehat{\lambda}_t \geq e^{-\delta}$ and farmers in

both countries use locally adapted inputs. In this special case, a sufficiently large market size makes type-S innovation privately profitable. The fixed cost of R&D is the same for both input types, while technological catch-up allows type-S innovation to earn a margin rate at least as high as that of frontier innovation. Given the maintained assumption that frontier innovation is profitable, R&D firms therefore choose to improve both input types.

With these conditions held over time, type-S inputs improve in every period, and the quality ratio λ_t rises toward the transferability ceiling χ due to technological catch-up. Southern farmers continue to use locally adapted technology, avoiding dynamic lock-in. This special case highlights how private innovation incentives depend on market size. When type-S inputs command a sufficiently large market, profit-maximizing firms improve them. In economies with smaller markets for type-S inputs, innovation can be unprofitable despite the scope for technological catch-up. The absence of innovation then widens the quality gap, which can further weaken demand and reinforce the initial disadvantage. Dynamic lock-in can thus emerge from privately optimal innovation and adoption decisions under laissez-faire.

Dynamic amplification effect. The two scenarios have starkly different implications for the evolution of the effective quality ratio $\mathcal{Q}(\widehat{\lambda}_t)$ and its contribution to agricultural productivity differences in (22). Along the dynamic lock-in path, a small market size makes type-S innovation unprofitable. As frontier type-N inputs improve, λ_t declines over time, eventually locking Southern farmers into inappropriate type-N inputs. The effective quality ratio is then $\mathcal{Q}(\widehat{\lambda}_t) = e^{-\delta}$, reflecting persistent environmental mismatch. In contrast, in the special case with a sufficiently large type-S market size, both input types improve each period. Technological catch-up raises λ_t toward the transferability ceiling χ, while Southern farmers continue to use locally adapted and improving type-S inputs. In this case, $\mathcal{Q}(\widehat{\lambda}_t) = \widehat{\lambda}_t$ converges to χ. Provided that $\chi > e^{-\delta}$, this limiting ratio is higher than $e^{-\delta}$. The contribution of effective input quality differences to agricultural productivity differences is therefore larger along the lock-in path, reflecting persistent environmental mismatch and forgone technological catch-up.

To quantify how much directed innovation amplifies agricultural productivity differences, I compare the decentralized equilibrium with a counterfactual in which both input types improve each period, holding economic fundamentals fixed across the two paths.[11] Given a common initial quality state $\lambda_{-1} \in (0, \chi]$, define the dynamic amplification factor as

$$\mathcal{A}_t := \frac{\text{APD}_t^U}{\text{APD}_t^D} = \left[\frac{\mathcal{Q}(\widehat{\lambda}_t^U)}{\mathcal{Q}(\widehat{\lambda}_t^D)}\right]^{\alpha/\Gamma}, \tag{23}$$

where APD_t^k denotes agricultural labor productivity differences $(Y_{S,t}^a/L_{S,t}^a)/(Y_{N,t}^a/L_{N,t}^a)$ evaluated along path $k \in \{D, U\}$. Here, D denotes the path under directed innovation, while U

[11] In the counterfactual, the entry condition (12) for type-S inputs is replaced in every period by the rule that type-S inputs are improved. All other equilibrium conditions remain in place, and the fixed cost of type-S innovation is incurred and borne by Northern households, who own the R&D firms.

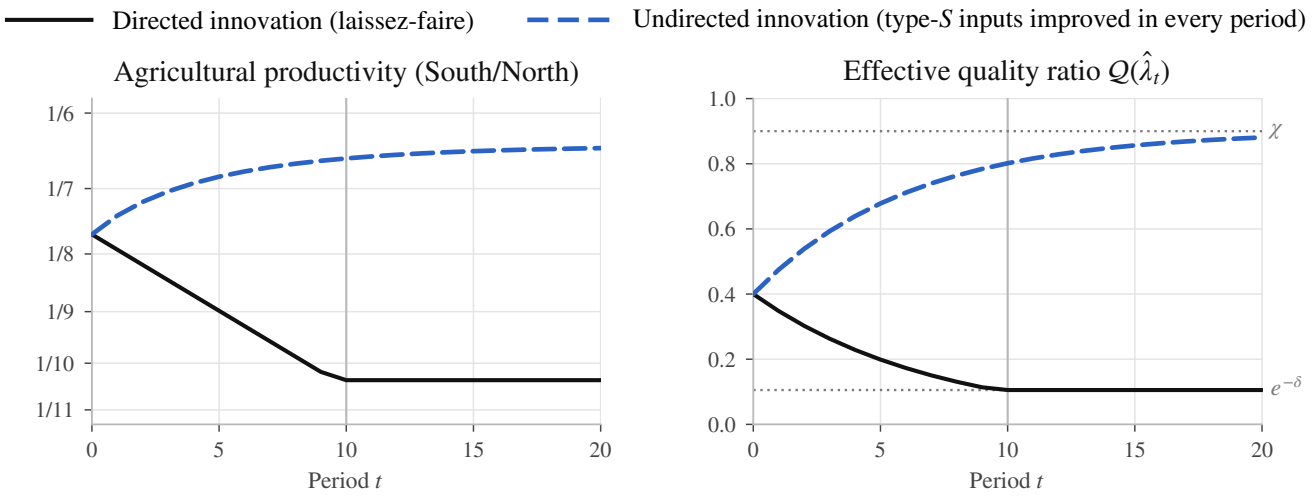


Figure 1: The dynamic lock-in and the counterfactual undirected innovation benchmark. The left panel shows the South-to-North agricultural labor productivity ratio. The right panel shows the effective quality ratio $\mathcal{Q}(\widehat{\lambda}_t)$. The solid line presents the laissez-faire path, and the dashed line presents the counterfactual path.

denotes the undirected innovation benchmark in which both input types improve each period. The factor exceeds one when Southern agricultural labor productivity relative to the North is lower under directed innovation than under the benchmark. Along the dynamic lock-in path without type-S innovation, the dynamic amplification factor $\mathcal{A}_t$ is nondecreasing over time and converges to $(\chi/e^{-\delta})^{\alpha/\Gamma}$. Equation (23) provides the paper's central measure of how much directed innovation amplifies agricultural labor productivity differences relative to the benchmark in which both input types improve each period.

An illustration. Figure 1 shows the two paths. In this example, the market for type-S inputs is too small for type-S innovation to be profitable, and both paths start from a quality ratio at which Southern farmers use type-S inputs. Along the laissez-faire path, R&D firms never improve type-S inputs, which fall further behind the frontier each period. Southern farmers continue to use them until period 10, when the price-adjusted quality ratio drops below $e^{-\delta}$ and they switch to inappropriate type-N inputs. From then on, the effective quality ratio stays at $e^{-\delta}$. The South never leaves this path under laissez-faire, because the quality gap only widens and the market for type-S inputs never reappears. Along the counterfactual path, the same economy from the same start has type-S inputs improved in every period. The quality ratio rises toward χ through technological catch-up, and Southern farmers continue to use locally adapted inputs, so the effective quality ratio rises with it. Since the first two sources of agricultural productivity differences in (22) are held fixed across the two paths, the paths differ only in the third, the effective quality ratio, which results from the interaction between innovation and adoption decisions. This interaction alone separates two long-run outcomes for the same econ-

omy: a permanently lower relative agricultural productivity with inappropriate technology, or a higher one with locally adapted technology that catches up with the frontier. The ratio of the two paths, the dynamic amplification factor (23), grows over time.

3 Constrained Efficiency

This section examines whether the dynamic lock-in described in Section 2 is efficient. I consider a social planner who takes the within-period equilibrium structure, including static monopoly distortions, as given and chooses only whether to improve type-S inputs. Section 3.1 formulates the constrained planner's problem, Section 3.2 characterizes constrained efficiency, and Section 3.3 explains why private and social innovation incentives differ and examines whether longer patents eliminate the difference.

3.1 The Constrained Planner's Problem

To isolate any inefficiency in the direction of innovation from the static monopoly distortion, I consider a social planner subject to three constraints. First, the planner faces the same technological and feasibility constraints as the decentralized economy. Second, the planner takes the within-period decentralized equilibrium structure as given. Third, the planner cannot redistribute income between Northern and Southern households, so each household consumes its own income, as under laissez-faire. Each period, the planner decides, on behalf of the R&D firms, whether to improve type-S inputs.[12] Since frontier technology improves in every period, this is the planner's only decision. The second restriction is central to this analysis. The static monopoly distortion remains uncorrected in both the decentralized equilibrium and the planner's allocation. The comparison therefore assesses whether private decisions about the direction of innovation are efficient, taking the within-period market structure as given.

Throughout this section, I assume that type-S innovation does not change Southern farmers' input choice in the period in which it occurs. They use type-S inputs in period t if and only if $\lambda_{t-1} \geq e^{-\delta}$.[13] Let $m_{S,t} \in \{0,1\}$ equal one if type-S inputs are improved in period t and zero otherwise.

Definition 2. Given $\lambda_{-1} \in (0,\chi]$, the constrained social planner chooses a sequence of type-S innovation decisions $\{m_{S,t}\}_{t=0}^{\infty}$ to maximize

$$\sum_{t=0}^{\infty} \frac{1}{(1+\rho)^t} \left[v_N u(C_{N,t}) + v_S u(C_{S,t}) \right], \qquad v_N, v_S \geq 0, \quad v_N + v_S = 1, \tag{24}$$

[12] This follows the constrained-planner approach of Stiglitz (1982) and Geanakoplos and Polemarchakis (1986), in which the planner makes a particular decision on behalf of private agents.

[13] When type-S innovation is non-drastic, this property holds without any further assumption. The assumption therefore matters only when innovation is drastic, and it holds at the parameter values calibrated in Section 5.

subject to (i) in every period t, the within-period equilibrium of Section 2.2 given $(\lambda_{t-1}, m_{S,t})$, which determines $(C_{N,t}, C_{S,t})$, and (ii) the law of motion (17) of the quality ratio.

Constraint (i) incorporates the three restrictions above. The technologies, endowments, and fixed costs of R&D are the same as in the decentralized economy. The planner replaces the type-S entry condition (12) with its innovation decision, while all other equilibrium conditions remain in place. The static monopoly distortion therefore remains uncorrected, and each household finances consumption from its own income without transfers between countries. In equation (24), v_N and v_S are the Pareto weights on the two households. I refer to the planner with $v_S = 1$ as the Southern planner, the planner with $v_N = 1$ as the Northern planner, and a planner with $v_N, v_S > 0$ as a global planner.

Recursive formulation. The planner's problem becomes stationary after removing the trend generated by frontier growth. Along any innovation path, consumption can be written as

$$C_{j,t} = (1+\xi)^{\alpha\phi t}\, \widetilde{C}_j(\lambda_{t-1}, m_{S,t}),$$

where $\widetilde{C}_j(\lambda, m_S)$ is detrended equilibrium consumption in country j given the inherited quality ratio λ and innovation decision m_S, with frontier quality normalized to one. The detrended planner's problem can be written recursively as

$$\begin{aligned} V(\lambda) &= \max_{m_S \in \{0,1\}} \Big\{ \mathcal{W}(\lambda, m_S) + \frac{(1+\xi)^{\alpha\phi(1-\sigma)}}{1+\rho} V\big(\lambda'(\lambda, m_S)\big) \Big\}, \\ \text{s.t.} \quad \lambda'(\lambda, m_S) &= \begin{cases} T(\lambda) & \text{if} \quad m_S = 1, \\ \lambda/(1+\xi) & \text{if} \quad m_S = 0. \end{cases} \end{aligned} \tag{25}$$

Here, V is the planner's value function, λ is the inherited quality ratio, and λ' is the quality ratio inherited in the next period. The period welfare function is $\mathcal{W}(\lambda, m_S) := v_N u(\widetilde{C}_N(\lambda, m_S)) + v_S u(\widetilde{C}_S(\lambda, m_S))$. It is evaluated at the within-period equilibrium specified in Definition 2 and thus incorporates the planner's within-period constraints.

For ease of exposition, I assume log utility ($\sigma = 1$) for the remainder of this section. Propositions 2, 3, and 5 below remain valid for general CRRA utility (see Appendix A.8).

The social and private values of type-S innovation. Define the social value of type-S innovation at the inherited quality ratio λ as

$$\Pi^{\text{SP}}(\lambda) := \Big[\mathcal{W}(\lambda, 1) + \frac{1}{1+\rho} V\big(T(\lambda)\big)\Big] - \Big[\mathcal{W}(\lambda, 0) + \frac{1}{1+\rho} V\Big(\frac{\lambda}{1+\xi}\Big)\Big]. \tag{26}$$

This is the welfare gain from improving type-S inputs in the current period rather than leaving them unchanged, with all subsequent innovation decisions chosen optimally in either case. The planner therefore improves type-S inputs if and only if $\Pi^{\text{SP}}(\lambda) \geq 0$. The corresponding private

value of type-S innovation is the R&D firm's net profit, $\Pi_{S,t}(\lambda_{t-1}) = \pi_{S,t}(\lambda_{t-1}) - c$. The firm improves type-S inputs if and only if $\Pi_{S,t}(\lambda_{t-1}) \geq 0$. Evaluating both decision rules at the same inherited quality ratio allows me to compare the planner's optimal innovation decision with the firm's.

Following Dávila and Korinek (2018), I call the decentralized equilibrium starting from λ_{-1} *constrained efficient* for Pareto weights (v_N, v_S) if its innovation sequence coincides with the sequence chosen by the planner with those weights from the same initial state, and *constrained inefficient* otherwise. Since the planner maximizes a weighted sum of the two households' lifetime welfare, a feasible innovation sequence that gives both households strictly higher lifetime welfare makes the equilibrium constrained inefficient for every choice of Pareto weights. I call this case *constrained Pareto inefficient*. Corollary 1 below provides sufficient conditions for this result from any initial quality ratio below $e^{-\delta}$.

3.2 The Private and the Social Value of Type-S Innovation

At a common inherited quality ratio $\lambda = \lambda_{t-1}$, the two decision rules can be written as

$$\text{laissez-faire:} \quad \Pi_{S,t}(\lambda) = \pi_{S,t}(\lambda) - c \geq 0,$$

$$\text{planner:} \quad \Pi^{\text{SP}}(\lambda) = \underbrace{\mathcal{W}(\lambda,1) - \mathcal{W}(\lambda,0)}_{\text{current-period effect}} + \underbrace{\frac{1}{1+\rho}\left[V\big(T(\lambda)\big) - V\Big(\frac{\lambda}{1+\xi}\Big)\right]}_{\text{continuation value}} \geq 0.$$

The private value is measured in numéraire units and the social value in welfare units. The private value includes only the firm's operating profit net of its R&D fixed cost. The social value has two components. The *current-period effect* captures all within-period welfare consequences of innovation: the R&D fixed cost, the profits accruing to households, and the resulting changes in agricultural revenues, prices, and farmers' profits. The *continuation value* is the difference in lifetime welfare between entering the next period with quality ratio $T(\lambda)$ and entering it with $\lambda/(1+\xi)$, with subsequent innovation decisions chosen optimally in either case. It captures the gains in later periods from the improved quality state that innovation leaves behind.

The first result identifies a key difference between the two decision rules: the planner accounts for the continuation value, whereas the R&D firm does not.

Proposition 2. *The private value of type-S innovation, $\Pi_{S,t}(\lambda_{t-1}) = \pi_{S,t}(\lambda_{t-1}) - c$, depends only on the current-period equilibrium and is independent of subsequent innovation decisions. The continuation value does not enter the firm's innovation decision.*

Proof. See Appendix A.9. □

This difference arises because the innovator has no claim on the future gains from the quality state that its innovation creates. The patent gives the firm the exclusive right to produce and sell improved inputs during the patent term, so the entry condition (12) compares the operating

profit of these inputs with the fixed cost of R&D. The innovation also raises the quality ratio that later periods inherit, and the gains from the higher ratio reach three groups: households in both countries, through the price of agricultural goods; the firms that improve type-S inputs in later periods, through their profits; and farmers, through their profits. The innovator has no claim on these gains that it can sell. Households receive them in later periods, and the firms that will build on the improved state have not yet entered. The intertemporal market through which the innovator could appropriate the gains is therefore missing, and Section 3.3 examines the three gains in turn. A longer patent extends the firm's claim on its own vintage, not to the improvements that follow. Proposition 5 below makes this precise.

When $\lambda_{t-1} < e^{-\delta}$, Southern farmers use type-N inputs, and the comparison becomes especially clear. Improved type-S inputs attract no demand in the current period, so innovation changes nothing in that period except that the fixed cost of R&D is incurred.

Proposition 3. *Fix $\lambda_{t-1} < e^{-\delta}$. The private value of type-S innovation is $\Pi_{S,t}(\lambda_{t-1}) = -c$. Its social value equals the continuation value net of the welfare loss from incurring the fixed cost of R&D. This welfare loss is the same for all $\lambda_{t-1} < e^{-\delta}$. The R&D firm therefore never improves type-S inputs, whereas the planner improves them if and only if the continuation value covers this welfare loss.*

Proof. See Appendix A.9. □

The firm's decision is straightforward. Under the assumption stated in Section 3.1, type-S innovation does not change Southern farmers' input choice within the period. At the price charged by the patent holder, the improved vintage offers too little quality per unit of expenditure to induce them to switch from type-N inputs. The firm therefore sells nothing and incurs the fixed cost.[14] The planner evaluates the welfare consequences of this cost, together with the continuation value. Below $e^{-\delta}$, the within-period equilibrium does not depend on the inherited quality ratio, so the welfare loss from incurring the cost is the same at every such ratio.

To illustrate when the gains from an additional innovation arise, compare two paths starting from the same inherited quality ratio $\lambda_{t-1} < e^{-\delta}$. Type-$S$ innovation occurs in period t on the first path but not the second, and in every subsequent period on both. Let $T^*(\lambda)$ denote the number of consecutive type-S innovations needed to raise the quality ratio from λ to at least $e^{-\delta}$. Only the first path incurs the fixed cost today, while Southern farmers initially use type-N inputs on both paths. For the next $T^*(\lambda_{t-1}) - 1$ periods, consumption for each household is identical across paths because Southern farmers continue using type-N inputs and both paths incur the same R&D cost. After $T^*(\lambda_{t-1})$ periods, the quality ratio on the first path reaches the adoption threshold, and Southern farmers switch to type-S inputs. On the second path, they continue using type-N inputs. Today's innovation therefore improves the quality state immediately, but its benefits emerge only when it brings forward type-S adoption. If continued

[14] Even if type-S innovation changed Southern farmers' input choice within the period, the Southern market might still be too small for the firm's operating profit to cover the fixed cost, as discussed in Section 2.3.2.

innovation is optimal on both paths, as under the condition in Proposition 4 (iii), the discounted welfare difference from period $t+1$ onward is the planner's continuation value. The planner then innovates when this value covers the current welfare cost of R&D.

Whether the continuation value covers the current welfare cost of R&D is a quantitative question. The cost is incurred today, while the gains arrive in later periods and are discounted. The social value of type-S innovation can be positive or negative. The following proposition therefore first characterizes constrained efficiency without restricting this sign, then considers the case in which the social value is positive at every quality ratio.

Proposition 4.

(i) Under laissez-faire, type-S innovation does not occur when $\lambda_{t-1} < e^{-\delta}$, or when $\lambda_{t-1} \geq e^{-\delta}$ and $\pi_{S,t}(\lambda_{t-1}) < c$.

(ii) If $\lambda_{-1} < e^{-\delta}$, type-S innovation never occurs under laissez-faire, and the quality ratio declines according to $\lambda_t = \lambda_{-1}/(1+\xi)^{t+1}$. If the social value is positive at some inherited quality ratio along this path, $\Pi^{SP}(\lambda_{-1}/(1+\xi)^t) > 0$ for some $t \geq 0$, the decentralized equilibrium is constrained inefficient, and the planner's innovation sequence yields strictly higher social welfare. If instead the social value is negative along the entire path, the planner chooses the laissez-faire innovation sequence, and the equilibrium is constrained efficient.

(iii) Suppose that $\Pi^{SP}(\lambda) > 0$ for all $\lambda \in (0, \chi]$. The planner then improves type-S inputs in every period, and the quality ratio converges to χ from any initial quality ratio. The decentralized equilibrium is constrained inefficient if $\Pi_{S,t} < 0$ in any period along the laissez-faire path, and constrained efficient if type-S innovation occurs in every period. In particular, it is constrained inefficient for every initial quality ratio $\lambda_{-1} < e^{-\delta}$.

Proof. See Appendix A.9. □

Corollary 1. *Suppose that the social value of type-S innovation is positive at every quality ratio for both the Southern and Northern planners. Improving type-S inputs in every period is then optimal for all Pareto weights (v_N, v_S). For any initial quality ratio $\lambda_{-1} < e^{-\delta}$, this path gives both households strictly higher lifetime welfare than the laissez-faire path. The decentralized equilibrium is therefore constrained Pareto inefficient.*

Proof. See Appendix A.9. □

Two remarks are in order. First, the inefficiency concerns dynamic lock-in. Starting from $\lambda_{-1} < e^{-\delta}$, Southern farmers use type-N inputs in every period under laissez-faire. Under the condition in Proposition 4 (iii), the planner instead improves type-S inputs in every period from the outset, and Southern farmers adopt them after $T^*(\lambda_{-1})$ periods. In the special case of Section 2.3.2 with symmetric economic fundamentals, a sufficiently large type-S market, and an initial quality ratio at or above $e^{-\delta}$, private firms improve type-S inputs in every period,

as does the planner under the condition in Proposition 4 (iii). In this case, the decentralized equilibrium is constrained efficient.

Second, Northern and Southern households bear different costs and receive different gains from type-S innovation. The Pareto weights determine how the planner values these effects. Northern households bear the fixed cost of R&D because they own the R&D firms. In later periods, improved inputs lower the agricultural price index at given agricultural revenues, benefiting households in both countries. Northern households also receive the profits of future type-S innovators, while the frontier patent holder loses its profits from sales to Southern farmers once they switch input types. Farmers' profits in each country follow its share of world agricultural revenue, which rises for the South and falls for the North as type-S inputs improve.

Southern households bear none of the fixed cost of R&D, whereas Northern households bear it in full and experience both gains and losses. The Northern planner's case thus makes the missing-market channel especially clear: even a planner concerned only with Northern households chooses type-S innovation where private firms do not, provided that their discounted future welfare gains cover the current welfare cost of R&D. This divergence between private and social decisions arises without assigning any welfare weight to Southern households.

3.3 Why the Private and the Social Value Differ

Proposition 2 shows that the private value of type-S innovation excludes the continuation value. This section examines what generates these future gains, who receives them, and why the original innovator cannot appropriate them. It then considers whether longer patents, or a persistent developer, would allow private firms to capture the gains.

Components of the continuation value. Under log utility, $\ln C_{j,t} = \ln E_{j,t} - \ln P_t$, where $E_{j,t}$ is household income in country j and P_t is the price index for the consumption aggregate. This index is proportional to $(P_t^a)^\phi$, where P_t^a is the agricultural price index.[15] In each subsequent period, the welfare difference between the path with type-S innovation in period t and the path without it is therefore

$$v_N \Delta \ln E_{N,t+s} + v_S \Delta \ln E_{S,t+s} - \phi \Delta \ln P_{t+s}^a, \qquad s \geq 1,$$

where Δ denotes the difference between the two paths, and the continuation value is the discounted sum of these welfare differences. The last term is the effect of agricultural prices, which is common to both households. The first two terms are differences in household income. Income consists of labor income, farmers' profits, and, in the North, the net profits of R&D firms. Labor income is the same along both paths, because the wage equals non-agricultural productivity in every use and the labor endowment is fixed. Income differences therefore arise through

[15] The price index for the consumption aggregate is given by $P_t = \phi^{-\phi}(1-\phi)^{-(1-\phi)}(P_t^a)^\phi$, and the agricultural price index is given by $P_t^a = \left[\mu (p_{N,t}^a)^{1-\varepsilon} + (1-\mu)(p_{S,t}^a)^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}}$.

farmers' profits and R&D profits. The continuation value thus reaches three groups: households in both countries, through agricultural prices; the R&D firms of later periods, through their profits; and farmers, through their profits.

Households, through agricultural prices. Once Southern farmers adopt type-S inputs, the path with the earlier innovation delivers higher effective input quality, which raises Southern agricultural productivity. At given agricultural revenues, this lowers the price of the Southern agricultural good, and hence the agricultural price index P^a. These gains reflect a knowledge spillover transmitted through prices. The knowledge embodied in the blueprint is nonrival, and competition passes its benefits to households in both countries. During the patent term, the quality improvement is fully reflected in the limit price, so the improved vintage offers the same quality per unit of expenditure as its predecessor. After the patent expires, the improved blueprint becomes available to the competitive fringe, which can supply the vintage at marginal cost. Households then benefit from lower agricultural prices without compensating the original innovator, who has no claim on these gains that it can sell.

Future innovators, through their profits. Earlier innovation gives future type-S innovators a better technological base. A higher inherited quality ratio can reduce the number of further innovations needed before Southern farmers adopt type-S inputs. After adoption, a higher quality ratio also increases the share of world agricultural revenue served by these inputs. The profits from subsequent innovations accrue to the firms that undertake them. The original innovator has no claim on these profits that it can sell, and future innovators have not yet entered. This creates a coordination problem across successive R&D firms. At inherited quality ratios below $e^{-\delta}$, each innovation is privately unprofitable. Yet a sequence of $T^*(\lambda)$ innovations raises the quality ratio to at least $e^{-\delta}$, allowing Southern farmers to adopt type-S inputs. If type-S innovation is profitable on $[e^{-\delta}, \chi]$, innovation continues under laissez-faire thereafter. Whether the sequence's operating profits cover its R&D costs in present value is a quantitative question. Even when they do, no individual firm can arrange the sequence, because it cannot contract with innovators in later periods.

Farmers, through their profits. As type-S inputs improve, the South's share of world agricultural revenue rises. Because demand for the Southern agricultural good is elastic ($\varepsilon > 1$), the increase in the quantity sold more than offsets the fall in its price. Farmers' profits are a constant share of agricultural revenue in each country, so their distribution shifts toward the South. The planner evaluates this redistribution according to the households' weighted marginal utilities of income and cannot offset it through transfers. These distributional effects enter the social value of innovation but do not, by themselves, establish inefficiency.

Together, these three channels account for the entire continuation value, none of which enters the original innovator's private return (see Proposition 2). The common problem is a missing intertemporal market. The innovator has no claim on these future gains that it can sell to

the households and firms that benefit from the improved quality state.

Longer patents. Under the condition of Proposition 4 (iii), the planner improves type-S inputs in every period. Starting from a quality ratio below $e^{-\delta}$, the first $T^*(\lambda_{-1})$ innovations produce vintages that attract no demand. Southern farmers then adopt type-S inputs, and innovation continues. The question is whether longer patents can bring the laissez-faire path closer to this path. A patent rewards innovation through sales of the patented vintage, and these sales occur only when Southern farmers use type-S inputs and no newer type-S vintage exists. Patent duration determines when the holder's exclusive rights over its vintage expire. The patent gives the holder no claim on the returns to subsequent innovations, even though these build on the improved quality state and generate the gains the planner values through eventual adoption and further catch-up. Extending patent duration can thus increase a vintage's return only if that vintage remains in demand and is not superseded during the additional period of protection.

The limitation is sharpest below $e^{-\delta}$. There, a longer patent protects a vintage for which there is no demand. Southern farmers use type-N inputs when the innovation occurs, and they do not adopt the improved vintage later, because its relative quality declines further as the frontier improves. Its operating profit is therefore zero throughout its patent term, regardless of patent duration, and the private value of the innovation remains $-c$. Adoption requires further innovations, and the original patent gives its holder no claim on the vintages they create.

The limitation also applies at or above $e^{-\delta}$, where Southern farmers use type-S inputs and a new vintage attracts demand. Suppose the market is too small for one period's operating profit to cover the fixed cost of R&D, so that type-S innovation does not occur under one-period patents. A longer patent could make an innovation profitable, but only if its vintage continued selling for several periods, which requires that no further type-S innovation occur in the meantime. Under a longer patent, the laissez-faire path therefore cannot coincide with the planner's path, on which type-S innovation occurs in every period.[16]

Proposition 5. *Suppose that patents last $T_p \geq 1$ periods, with blueprints disclosed at expiry.*

(i) For any $\lambda_{t-1} < e^{-\delta}$, the type-S vintage created in period t attracts no Southern demand during its patent term, regardless of subsequent innovation decisions. The private value of type-S innovation is therefore $-c$ for every T_p, and no type-S innovation occurs under laissez-faire below $e^{-\delta}$.

(ii) A type-S patent earns operating profit only in periods in which Southern farmers use type-S inputs and its vintage is the newest type-S vintage. If another type-S innovation follows in the next period, the private value of innovation in period t is therefore $\pi_{S,t}(\lambda_{t-1}) - c$ for every T_p. Hence

[16] A patent holder considering another innovation must also account for the profits it would forgo by replacing its existing vintage, Arrow (1962)'s replacement effect. A firm that undertook every improvement would still earn at most one period's operating profit from each vintage along this path. Its additional incentive would come from its claim on subsequent innovations, as in the persistent developer considered below.

type-S innovation occurs in every period under laissez-faire only if it does so under one-period patents.

Proof. See Appendix A.10. □

Part (ii) applies at every quality ratio and market size. It follows from the displacement of each vintage by the next improvement within an input type and does not require the baseline's bang-bang input choice. At or above $e^{-\delta}$, a longer patent can induce an innovation that would not occur under one-period patents and thereby bring the laissez-faire path closer to the planner's path. Each such innovation, however, must be followed by a period without type-S innovation, so longer patents cannot sustain innovation in every period where one-period patents do not. They also induce no innovation below $e^{-\delta}$. Longer patents therefore do not eliminate the difference between the two paths.

A persistent developer. A persistent developer that holds the type-S research opportunities in every period can internalize profits across successive innovations. It chooses improvements to maximize its discounted net profits. In practice, such a developer might hold exclusive rights to develop varieties for a crop in a region or control the germplasm on which subsequent varieties depend. A public research institute, such as Embrapa in Brazil (Akerman et al., 2026), can also sustain research across periods while pursuing public objectives. Public funding allows it to support improvements that private firms would forgo, as does an R&D subsidy.

The developer internalizes profits across the whole sequence of type-S innovations. Profits after adoption can therefore make it worthwhile to bear the costs of the first $T^*(\lambda_{-1})$ innovations, whose vintages attract no demand. If it improves type-S inputs in every period, each vintage earns operating profit for at most one period, regardless of patent duration. The developer's additional incentive therefore comes from its claim on the returns to subsequent innovations. In this sense, it accounts for more of the planner's continuation value than a longer patent does. Of the three components, however, it internalizes only the profits from future type-S innovation. It does not count households' gains from lower agricultural prices or farmers' profits.

Summary. Under the condition in Proposition 4 (iii), private innovation decisions differ from the planner's wherever type-S innovation is unprofitable under laissez-faire, including below $e^{-\delta}$. Below this threshold, the disagreement reflects the continuation value: the firm earns no operating profit, while the planner counts the future gains from the improved quality state. The firm can foresee these gains but has no claim on them that it can sell. An R&D subsidy can close this gap through a payment conditional on innovation (see Appendix A.10). A longer patent does not close it below the adoption threshold, because it protects a vintage that never sells and provides no claim on subsequent innovations. A persistent developer internalizes profits across the sequence but still excludes the gains to households and farmers that the planner counts.

4 Quantitative Model

This section extends the baseline model of Section 2 for quantitative analysis. Three extensions relax its simplifying assumptions. First, intermediate inputs comprise a continuum of product lines with different R&D fixed costs, so R&D firms choose not only whether to improve an input type but also how many of its product lines to improve. Innovation intensity, the share of product lines improved, replaces the binary innovation decision and introduces an intensive margin of directed innovation. Second, farmers differ in how suitable each input type is for their land. Adoption shares are therefore no longer binary: both input types are used in each country, and their shares vary continuously with the price-adjusted quality ratio. Third, I introduce weak enforcement of intellectual property rights in the South and an additional R&D fixed cost for type-S inputs relative to type-N inputs. Weak enforcement allows local imitators to capture part of innovators' operating profits from sales to Southern farmers. The additional cost reflects Northern R&D firms' limited knowledge of Southern agronomic conditions and the field trials needed for adaptation. Both features weaken incentives to improve type-S inputs.

Four additional features help the model match the data: a labor wedge that keeps agricultural wages below non-agricultural wages; a subsistence requirement for agricultural consumption; iceberg trade costs; and an explicit land endowment. Section 4.1 describes the environment, Section 4.2 defines the equilibrium and derives the law of motion for the quality ratio, and Section 4.3 examines agricultural productivity differences and revisits the key mechanism.

4.1 Environment

The basic setting is as in Section 2.1, with two changes to trade and labor mobility. First, agricultural goods are subject to iceberg trade costs between countries, while the non-agricultural good remains freely traded. Second, a labor wedge keeps agricultural wages below non-agricultural wages within each country.

4.1.1 Final Good Producers

Agriculture. Each country j is populated by a continuum of farmers indexed by $i \in [0,1]$, each farming Z_j units of land. Intermediate inputs comprise a continuum of product lines indexed by $\ell \in [0,1]$. Each line represents a distinct input category, such as an improved seed for a given crop, a fungicide, or a nitrogen fertilizer, and has both North-adapted and South-adapted versions. Farmers differ in how suitable each input type is for their land. Farmer i in country j employs a constant-returns-to-scale agricultural production function:

$$Y^a_{j,t}(i) = A^a_j \left(L^a_{j,t}(i)\right)^{\beta} Z_j^{\gamma} \int_0^1 \Big[\sum_{k \in \{N,S\}} v_{j,k,t}(i)\, B_{j,k}\, q_{k,\ell,t}\, x_{j,k,\ell,t}(i) \Big]^{\alpha} \mathrm{d}\ell, \tag{27}$$

where $Y^a_{j,t}(i)$ and $L^a_{j,t}(i)$ denote agricultural output and labor of farmer i, $x_{j,k,\ell,t}(i)$ denotes the quantity of product line ℓ of type-k inputs used by farmer i, $q_{k,\ell,t}$ denotes the base quality of

product line ℓ of type-k inputs, and $v_{j,k,t}(i)$ captures the suitability of type-k inputs for the land of farmer i. The agricultural production function now explicitly includes land, with share $\gamma = 1 - \alpha - \beta$, and thus exhibits constant returns to scale in land, labor, and intermediate inputs.

Within each product line, the two input types are perfect substitutes in effective units, while inputs from different product lines are imperfect substitutes. The effective quality of an input now has three components: its base quality q, its appropriateness B for the country's agricultural environment, and its suitability v for the land of the farmer who uses it. Suitability is drawn independently across farmers, input types, and periods from a Fréchet distribution:

$$\Pr\{v_{j,k,t}(i) \leq x\} = \exp(-x^{-\theta}),$$

where θ governs the dispersion of suitability across farmers: a lower θ corresponds to a greater dispersion.[17] Suitability captures how well conditions on a farmer's land, such as soil and rainfall, match the environment for which type-k inputs are designed. For example, type-N inputs are more effective for Southern farmers whose land conditions more closely resemble those in the North.

Farmer i in country j chooses labor $L^a_{j,t}(i)$ and the quantities $\{x_{j,k,\ell,t}(i)\}_{k\in\{N,S\},\,\ell\in[0,1]}$ of both input types on every product line to maximize profits,

$$p^a_{j,t}\, Y^a_{j,t}(i) - w^a_{j,t}\, L^a_{j,t}(i) - \int_0^1 \Big[\sum_{k\in\{N,S\}} p^x_{k,\ell,t}\, x_{j,k,\ell,t}(i) \Big] \mathrm{d}\ell, \tag{28}$$

subject to the agricultural production function (27), taking as given the agricultural output price $p^a_{j,t}$, the agricultural wage $w^a_{j,t}$, and the input prices $p^x_{k,\ell,t}$ for $k \in \{N,S\}$ and $\ell \in [0,1]$. Because the two input types are perfect substitutes within each product line, the farmer chooses, for each product line, the type with the higher effective quality per unit of expenditure, $v_{j,k,t}(i)\, B_{j,k}\, q_{k,\ell,t} / p^x_{k,\ell,t}$. When type-$S$ innovation is non-drastic, all product lines of a given input type offer the same quality per unit of expenditure, whether or not they are improved, as explained in Section 4.1.2. The comparison therefore has the same outcome on every product line, and each farmer uses the same input type across all of them.[18]

Given the input type, the farmer allocates expenditure across its product lines. Define the price-adjusted quality of type-k inputs and their price-adjusted effective quality in country j as

$$\widetilde{q}_{k,t} := \left[\int_0^1 \left(\frac{q_{k,\ell,t}}{p^x_{k,\ell,t}} \right)^{\frac{\alpha}{1-\alpha}} \mathrm{d}\ell \right]^{\frac{1-\alpha}{\alpha}} \qquad \text{and} \qquad h_{j,k,t} := B_{j,k}\, \widetilde{q}_{k,t}. \tag{29}$$

[17] I assume $\theta > \alpha/\gamma$, so that aggregate agricultural revenue is finite.

[18] When type-S innovation is drastic and only some type-S product lines are improved, the improved lines offer higher quality per unit of expenditure than the unimproved ones. A farmer may then prefer type-S inputs on the improved lines and type-N inputs on the others. In this case, I assume each farmer commits to one input type across all product lines. This assumption captures the idea that farmers choose their inputs as a package before the season begins. This case does not arise in the calibrated model along any path examined in Sections 5 and 6. See Appendix B.3 for details.

The price-adjusted quality $\widetilde{q}_{k,t}$ measures the quality of type-k inputs per unit of expenditure when expenditure is allocated optimally across product lines. With a single product line, $\widetilde{q}_{k,t} = q_{k,t}/p^x_{k,t}$, and $h_{j,k,t}$ reduces to its baseline expression in (5). Farmer i's optimal input type satisfies $\tau^*_{j,t}(i) \in \arg\max_{k\in\{N,S\}} v_{j,k,t}(i)\, h_{j,k,t}$, which generalizes the baseline input choice rule in (4). Because suitability varies across farmers, both input types are used in each country. The share of farmers in country j using type-k inputs is

$$\omega_{j,k,t} = \frac{h^\theta_{j,k,t}}{h^\theta_{j,N,t} + h^\theta_{j,S,t}}, \tag{30}$$

which also equals the share of country j's total intermediate input expenditure allocated to type-k inputs. Unlike in the baseline model, adoption shares vary smoothly with the price-adjusted effective qualities of the two input types. If the price-adjusted quality of type-S inputs falls relative to that of type-N inputs, fewer farmers in each country use type-S inputs, but some still use them. As $\theta \to \infty$, the adoption shares converge to those implied by the bang-bang input choice in Section 2.1.1.

As in the baseline model, farmer i's optimal labor demand is $L^a_{j,t}(i) = \beta\, p^a_{j,t} Y^a_{j,t}(i)/w^a_{j,t}$. Define aggregate agricultural output and labor in country j as $Y^a_{j,t} \coloneqq \int_0^1 Y^a_{j,t}(i)\,\mathrm{d}i$ and $L^a_{j,t} \coloneqq \int_0^1 L^a_{j,t}(i)\,\mathrm{d}i$, respectively.

Non-agriculture. The non-agricultural sector is as in Section 2.1.1: the representative firm produces $Y^n_{j,t} = A^n_j L^n_{j,t}$, and its optimality condition is $w^n_{j,t} = A^n_j$.

4.1.2 R&D Firms and Intermediate Input Producers

Innovation now occurs at the product-line level. Research on a product line of type-k inputs raises its base quality as in the baseline model. Three features of the R&D block are new: knowledge from improved product lines is shared across lines of the same input type when patents expire; weak enforcement of intellectual property rights in the South allows local imitators to capture part of patent holders' operating profits from sales to Southern farmers; and R&D fixed costs vary across product lines and are higher on average for type-S inputs.

Timing. In each period, one R&D firm per input type, located in the North, holds the research opportunities for all product lines of that input type. Each product line has a manager who observes its R&D fixed cost and decides whether to improve it. Managers do not coordinate their decisions and take aggregate variables as given.[19] The three stages of Section 2.1.2 apply to each product line. A manager who improves a product line in Stage 1 sells the improved inputs as a patent holder in Stage 2, and the patent expires in Stage 3.

[19] The assumption that one firm holds all research opportunities for an input type is not essential. The same product lines would be improved if each were open to the pool of potential R&D firms in Section 2.1.2, with patent exclusion deterring a second firm from improving the same line.

There is one addition to Stage 3. When patents expire, the knowledge embodied in the improved product lines becomes available to all product lines of the same input type. The common base quality of type-k inputs is the geometric mean of the product-line base qualities:

$$\ln q_{k,t} = \int_0^1 \ln q_{k,\ell,t}\, \mathrm{d}\ell. \tag{31}$$

Every product line of type-k inputs begins period $t+1$ at base quality $q_{k,t}$. The competitive fringe can produce inputs at this quality, and research in period $t+1$ builds on it. This assumption captures the idea that knowledge of local soils, climate, and pests is useful across input categories when adapting technologies to the same agricultural environment. It also allows the base qualities of the two input types to summarize the state of technology.

Stage 2: Production. As in the baseline model, producing one unit of any product line requires ψ units of the non-agricultural good. Product lines not improved in period t are supplied by a competitive fringe at marginal cost. For an improved product line ℓ of type-k inputs, the patent holder's optimal price is

$$p^x_{k,\ell,t} = \min\{\alpha,\ G_{k,\ell,t}\,\psi\},$$

where $G_{k,\ell,t} \coloneqq q_{k,\ell,t}/q_{k,t-1}$ is the gross quality step relative to the common base quality $q_{k,t-1}$ inherited from period $t-1$. All improved product lines of a given input type have the same gross step size, denoted by $G_{k,t}$.

Intellectual property rights are weakly enforced in many developing countries. To capture this, I assume that imitators in the South appropriate a fraction $1-\zeta_S$ of a patent holder's operating profit from sales to Southern farmers, where $\zeta_S \in (0,1]$ measures the strength of enforcement of intellectual property rights in the South. These imitators are owned by Southern households. The patent holder therefore retains all operating profit from sales to Northern farmers and a fraction ζ_S of that from sales to Southern farmers.

When type-k innovation is non-drastic, all product lines of that input type offer the same quality per unit of expenditure, so farmers spend equally across them. The patent holder's operating profit from an improved product line of type-k inputs is then

$$\pi_{k,t} = \nu(G_{k,t})\left[\omega_{N,k,t}\cdot\alpha\, p^a_{N,t} Y^a_{N,t} + \zeta_S\,\omega_{S,k,t}\cdot\alpha\, p^a_{S,t} Y^a_{S,t}\right], \tag{32}$$

where $\nu(\cdot)$ is the margin rate in (9). As in equation (8), operating profit equals the margin rate times market size. There are three differences from the baseline model. First, adoption shares $\omega_{j,k,t}$ lie strictly between zero and one, so both input types have a market in both countries. Second, the patent holder retains only a fraction ζ_S of the operating profit from its Southern market. Third, $\pi_{k,t}$ is operating profit per improved product line, common to all improved lines of a given input type.[20]

[20] When type-S innovation is drastic, farmers spend more on improved type-S product lines than on unimproved ones. The profit expression then requires an additional factor, given in Appendix B.4.

Stage 1: Entry and R&D decisions. The R&D block specifies quality improvements and R&D fixed costs. At the beginning of period t, every product line of type-k inputs inherits the base quality $q_{k,t-1}$. If product line ℓ is improved in period t, its base quality becomes

$$q_{N,\ell,t} = (1+\xi)\, q_{N,t-1} \qquad \text{or} \qquad q_{S,\ell,t} = (1+\xi)\left[a\,\chi\, q_{N,t-1} + (1-a)\, q_{S,t-1}\right], \tag{33}$$

for type-N and type-S inputs, respectively. These are the same quality improvements as in the baseline model, now applied to individual product lines: type-N innovation raises quality by the factor $1+\xi$, while type-S innovation also incorporates catch-up as in (10). If product line ℓ is not improved, $q_{k,\ell,t} = q_{k,t-1}$. All improved product lines of a given input type therefore have the same gross step size: $G_{N,t} = 1+\xi$ and $G_{S,t} = G_S(\lambda_{t-1})$, where G_S is defined in (18).

R&D fixed costs differ across product lines because the number of researchers required to improve them varies. Let $c_{k,\ell,t}$ denote the R&D fixed cost of improving product line ℓ of type-k inputs in period t. Costs are drawn independently across product lines, input types, and periods. For type-k inputs, the cost distribution has mean c_k, support $[\underline{c}_k, \overline{c}_k]$ with $\underline{c}_k > 0$, and continuous distribution function $H(c/c_k)$. The mean R&D fixed costs are

$$c_N \coloneqq \frac{w^r_{N,t}}{\eta} \qquad \text{and} \qquad c_S \coloneqq (1+f)\, c_N, \tag{34}$$

where c_N is the baseline R&D fixed cost in (11) and $f \geq 0$ captures the additional cost of improving type-S inputs. Northern R&D firms are less familiar with Southern agricultural environments than with Northern ones, so adapting inputs to Southern conditions requires additional research, such as field trials in the South.

The manager of product line ℓ of type-k inputs improves it if and only if anticipated operating profit covers the R&D fixed cost, $\pi_{k,t} \geq c_{k,\ell,t}$. This is the baseline entry condition (12) applied to a product line. Because anticipated operating profit is common to all product lines of a given input type, those with the lowest R&D fixed costs are improved. I define the innovation intensity of type-k inputs as the share of its product lines that are improved in period t:

$$m_{k,t} \coloneqq \int_0^1 \mathbb{1}\left\{\pi_{k,t} \geq c_{k,\ell,t}\right\} \mathrm{d}\ell = H\left(\frac{\pi_{k,t}}{c_k}\right) \in [0,1]. \tag{35}$$

Innovation intensity captures two margins of directed innovation. The extensive margin is whether type-k innovation occurs at all: $m_{k,t} > 0$ if and only if $\pi_{k,t} > \underline{c}_k$. The intensive margin is how many product lines are improved: above $\underline{c}_k$, higher operating profit raises innovation intensity until all product lines are improved. The baseline model has only the extensive margin. If all product lines had the same R&D fixed cost, the entry condition (12) would imply an innovation intensity of either zero or one. For type-S inputs, this would coincide with the binary innovation decision $m_{S,t} \in \{0,1\}$.

Finally, the R&D firm's net profit from type-k inputs is the sum of net profits across its improved product lines:

$$\Pi_{k,t} \coloneqq \int_0^1 \mathbb{1}\left\{\pi_{k,t} \geq c_{k,\ell,t}\right\}\left(\pi_{k,t} - c_{k,\ell,t}\right) \mathrm{d}\ell.$$

4.1.3 Households and Market Clearing

Households. The representative household in country j has preferences given by (13). The consumption aggregate is now given by

$$C_{j,t} = \left(C^a_{j,t} - \bar{a}\, L_j\right)^{\phi} \left(C^n_{j,t}\right)^{1-\phi}, \tag{36}$$

where $\bar{a} \geq 0$ is the subsistence requirement for agricultural goods per household member, and $C^a_{j,t}$ is the CES aggregate in (14), with equal weights on the two agricultural varieties. When $\bar{a} > 0$, the expenditure share of agricultural goods declines with income. This income effect is the source of structural change from agriculture to non-agriculture in this model.[21]

Trade in agricultural goods is subject to iceberg trade costs, as in Tombe (2015): delivering one unit of an agricultural good to the other country requires shipping $\tau_g \geq 1$ units. These costs help the model match the high share of domestically produced goods in agricultural consumption. Let

$$\tau_{j,k} := \tau_g^{\mathbb{1}\{k \neq j\}} = \begin{cases} 1 & \text{if } k = j, \\ \tau_g & \text{if } k \neq j, \end{cases}$$

so that the household in country j pays $\tau_{j,k}\, p^a_{k,t}$ per unit of the agricultural variety produced in country k. The household's period budget constraint is

$$\sum_{k \in \{N,S\}} \tau_{j,k}\, p^a_{k,t}\, c^a_{j,k,t} + C^n_{j,t} = E_{j,t}, \tag{37}$$

where $E_{j,t}$ is household income, defined below. The agricultural price index in country j is

$$P^a_{j,t} := \left[\mu \left(\tau_{j,N}\, p^a_{N,t}\right)^{1-\varepsilon} + (1-\mu)\left(\tau_{j,S}\, p^a_{S,t}\right)^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}}. \tag{38}$$

The household's optimal expenditure on agricultural goods is

$$E^a_{j,t} := \sum_{k \in \{N,S\}} \tau_{j,k}\, p^a_{k,t}\, c^a_{j,k,t} = \phi\, E_{j,t} + (1-\phi)\, P^a_{j,t}\, \bar{a}\, L_j. \tag{39}$$

The household first covers the cost of subsistence consumption and then spends a share ϕ of its remaining income on agricultural goods. When $\bar{a} = 0$, this reduces to the baseline expenditure rule, $E^a_{j,t} = \phi E_{j,t}$.

Household income is given by

$$\begin{aligned} E_{N,t} &= w^a_{N,t} L^a_{N,t} + w^n_{N,t} L^n_{N,t} + w^r_{N,t} L^r_{N,t} + \text{Profit}^a_{N,t} + \sum_{k \in \{N,S\}} \Pi_{k,t}, \\ E_{S,t} &= w^a_{S,t} L^a_{S,t} + w^n_{S,t} L^n_{S,t} + \text{Profit}^a_{S,t} + \text{Profit}^x_{S,t}, \end{aligned} \tag{40}$$

where $\text{Profit}^a_{j,t}$ is the profits of farmers in country j, the return to the land they own, $\Pi_{k,t}$ is the net profit of the R&D firm for type-k inputs, and $\text{Profit}^x_{S,t}$ is Southern imitators' profits. As in the baseline model, households receive labor income and the profits of the firms they own.

[21] See, e.g., Herrendorf, Rogerson, and Valentinyi (2013), Boppart (2014), Comin, Lashkari, and Mestieri (2021).

Labor wedge. Following Restuccia et al. (2008) and Tombe (2015), I introduce a wedge between agricultural and non-agricultural wages:

$$w^a_{j,t} = \frac{w^n_{j,t}}{1+\tau_{w,j}}, \tag{41}$$

where $\tau_{w,j} \geq 0$ is the labor wedge in country j. Workers in agriculture are paid a fraction $1/(1+\tau_{w,j})$ of the non-agricultural wage. The wedge captures the observed wage gap between agriculture and non-agriculture. Researchers earn the non-agricultural wage, $w^r_{N,t} = w^n_{N,t}$.

Market clearing. Market clearing for agricultural goods requires

$$\tau_{N,j}\, c^a_{N,j,t} + \tau_{S,j}\, c^a_{S,j,t} = Y^a_{j,t}, \qquad \forall j \in \{N,S\},$$

where the left-hand side includes goods lost in transit. For the non-agricultural good, market clearing requires

$$C^n_{N,t} + C^n_{S,t} + \sum_{j\in\{N,S\}} \sum_{k\in\{N,S\}} \int_0^1 \int_0^1 \psi\, x_{j,k,\ell,t}(i)\, \mathrm{d}i\, \mathrm{d}\ell = Y^n_{N,t} + Y^n_{S,t}.$$

The labor market clearing conditions are as in Section 2.1.3. The number of Northern researchers now depends on how many product lines are improved. Each product line's fixed cost of R&D equals the wage bill for its researchers, so $w^r_{N,t} L^r_{N,t} = \sum_k \int_0^1 \mathbb{1}\{\pi_{k,t} \geq c_{k,\ell,t}\}\, c_{k,\ell,t}\, \mathrm{d}\ell$.

4.2 Equilibrium

The definition of equilibrium is as follows.

Definition 3. Given $(q_{N,-1}, q_{S,-1})$ with $\lambda_{-1} \in (0,\chi]$, an equilibrium consists of a sequence of prices $\left\{ \left(p^a_{j,t}, w^a_{j,t}, w^n_{j,t}\right)_{j\in\{N,S\}}, w^r_{N,t}, \left(p^x_{k,\ell,t}\right)_{k\in\{N,S\},\ell\in[0,1]} \right\}_{t=0}^{\infty}$ and corresponding quantities such that, in every period t, (i) farmers maximize (28) subject to (27), taking prices as given; (ii) non-agricultural final good producers maximize (6), taking prices as given; (iii) the patent holder of each improved product line maximizes its operating profit subject to farmers' demand, and product lines that are not improved are supplied by the competitive fringe at marginal cost; (iv) the manager of product line ℓ of type-k inputs improves it if and only if $\pi_{k,t} \geq c_{k,\ell,t}$; (v) households maximize (13) with the consumption aggregate (36), subject to (37) and (40); (vi) the base qualities of product lines evolve according to (33), and the base qualities of input types are given by (31); (vii) wages satisfy (41) and $w^r_{N,t} = w^n_{N,t}$; and (viii) the markets for the two agricultural goods, the non-agricultural good, and labor clear.

Recursive structure of the equilibrium. The equilibrium remains recursive, but the inherited quality ratio λ_{t-1} is no longer the only state variable. Subsistence consumption is fixed in

units of agricultural goods, while agricultural output grows with input quality. The period-t equilibrium therefore also depends on the level of frontier quality through $\bar{a}\, q_{N,t-1}^{-\alpha}$.

Law of motion of the quality ratio. The law of motion for the quality ratio $\lambda_t = q_{S,t}/q_{N,t}$ now depends on innovation intensities. By (31) and (33), the base quality of type-k inputs evolves according to

$$q_{k,t} = \left(G_{k,t}\right)^{m_{k,t}} q_{k,t-1}.$$

Base quality therefore grows by the gross quality step raised to the share of product lines improved. Taking the ratio of the expressions for type-S and type-N inputs yields

$$\lambda_t = \lambda_{t-1} \frac{G_S(\lambda_{t-1})^{m_{S,t}}}{(1+\xi)^{m_{N,t}}}.$$

As in Section 2.2, the main focus is on type-S innovation. To focus on this margin, I assume that parameters are such that $\pi_{N,t} \geq \bar{c}_N$ for all t, so all type-N product lines are improved in every period ($m_{N,t} = 1$). The law of motion for λ_t then becomes

$$\lambda_t = T(\lambda_{t-1})^{m_{S,t}} \left(\frac{\lambda_{t-1}}{1+\xi}\right)^{1-m_{S,t}}, \tag{42}$$

where $T(\lambda_{t-1}) = (1-a)\lambda_{t-1} + a\chi$ as in (17). Equation (42) expresses the new quality ratio as a weighted geometric mean of the two branches of the baseline law of motion (17), with weights given by the shares of improved and unimproved type-S product lines. If all type-S product lines are improved ($m_{S,t} = 1$), the quality ratio becomes $T(\lambda_{t-1})$. If none is improved ($m_{S,t} = 0$), it falls to $\lambda_{t-1}/(1+\xi)$. For intermediate innovation intensities, the quality ratio lies between these two values.

The innovation intensity required to keep the quality ratio constant is

$$\overline{m}(\lambda_{t-1}) := \frac{\ln(1+\xi)}{\ln G_S(\lambda_{t-1})}. \tag{43}$$

By (42), $\lambda_t > \lambda_{t-1}$ if and only if $m_{S,t} > \overline{m}(\lambda_{t-1})$. There are three cases. First, if $m_{S,t} = 0$, no type-S product lines are improved, and the quality ratio falls by the factor $1/(1+\xi)$. Second, if $0 < m_{S,t} < \overline{m}(\lambda_{t-1})$, the base quality of type-S inputs rises, but more slowly than that of type-N inputs, so the quality ratio still falls. This case does not arise in the baseline model, where type-S innovation raises the quality ratio whenever $\lambda_{t-1} < \chi$. Third, if $m_{S,t} > \overline{m}(\lambda_{t-1})$, the quality ratio rises. The threshold $\overline{m}(\lambda_{t-1})$ is strictly increasing in λ_{t-1} and equals one at $\lambda_{t-1} = \chi$. When type-S inputs lag far behind the frontier, catch-up generates large proportional quality gains on improved product lines, so improving a small share of lines suffices to keep pace with the frontier. As the quality ratio approaches χ, the quality step on each improved line approaches the frontier step, and a larger share of lines must be improved to keep pace.

In equilibrium, the innovation intensity of type-S inputs is determined by the entry decisions of managers:

$$m_{S,t} = H\left(\frac{\pi_{S,t}(\lambda_{t-1})}{c_S}\right).$$

Operating profit $\pi_{S,t}(\lambda_{t-1})$ depends on the inherited quality ratio through adoption shares and the margin rate. This entry condition and (42) therefore jointly determine the evolution of the quality ratio.[22] As in the baseline model, the quality ratio converges to χ if all type-S product lines are improved in every period, and to zero if none is ever improved.

Drastic and non-drastic innovation. All improved product lines of type-S inputs share the gross step size $G_S(\lambda_{t-1})$, so Lemma 1 applies to each. Type-S innovation is therefore drastic if and only if $\lambda_{t-1} \leq \lambda^{\mathrm{D}}$, and each improved line has the margin rate $\nu_{S,t}$ given in (20).

Allocation-relevant quality ratio. As in Section 2.2, farmers' input choices and the adoption shares in (30) depend on the price-adjusted quality ratio. With multiple product lines, this ratio is defined as $\widehat{\lambda}_t := \widetilde{q}_{S,t}/\widetilde{q}_{N,t}$ and is given by

$$\widehat{\lambda}_t = \begin{cases} \lambda_{t-1} & \text{if } \lambda_{t-1} > \lambda^{\mathrm{D}}, \\ \left[m_{S,t} \left(\alpha(1+\xi)\, T(\lambda_{t-1}) \right)^{\frac{\alpha}{1-\alpha}} + (1 - m_{S,t})\, \lambda_{t-1}^{\frac{\alpha}{1-\alpha}} \right]^{\frac{1-\alpha}{\alpha}} & \text{if } \lambda_{t-1} \leq \lambda^{\mathrm{D}}. \end{cases} \tag{44}$$

When type-S innovation is non-drastic, the quality improvement is fully reflected in the patent holder's price, so improved and unimproved product lines offer the same quality per unit of expenditure. The price-adjusted quality ratio therefore equals the inherited quality ratio regardless of innovation intensity. Current innovation then affects adoption shares only from period $t+1$ onward. When type-S innovation is drastic, improved product lines offer higher quality per unit of expenditure, allowing part of the quality gain to reach farmers within the period. This effect is larger when more product lines are improved. At $m_{S,t} = 1$ and $m_{S,t} = 0$, (44) reduces to the baseline expressions in (21) with and without type-S innovation, respectively.

4.3 Equilibrium Characterization and Key Mechanism

This section examines agricultural productivity differences in the quantitative model and revisits the key mechanism.

4.3.1 Agricultural Productivity Differences

The following proposition expresses agricultural labor productivity differences in terms of the effective quality of intermediate inputs and agricultural employment in the two countries.

[22] Operating profit also depends on innovation intensity through household incomes and, when type-S innovation is drastic, through adoption shares and farmers' expenditure on each improved product line. Multiple innovation intensities may therefore satisfy the entry condition, in which case I select the lowest. Appendix B.4 discusses equilibrium existence, and Appendix C.9 verifies uniqueness at the calibrated parameters.

Proposition 6. *In equilibrium, agricultural labor productivity differences between the South and the North in period t are given by*

$$\frac{Y^a_{S,t}/L^a_{S,t}}{Y^a_{N,t}/L^a_{N,t}} = \left(\frac{A^a_S}{A^a_N}\right) \cdot \left(\frac{Z_S}{Z_N}\right)^{\gamma} \cdot \left(\frac{A^n_S}{A^n_N} \cdot \frac{1+\tau_{w,N}}{1+\tau_{w,S}}\right)^{\alpha} \cdot \widetilde{\mathcal{Q}}(\widehat{\lambda}_t)^{\alpha} \cdot \left(\frac{L^a_{S,t}}{L^a_{N,t}}\right)^{-\gamma}, \tag{45}$$

where

$$\widetilde{\mathcal{Q}}(\widehat{\lambda}) := \left[\frac{e^{-\delta\theta} + \widehat{\lambda}^{\theta}}{1 + \left(e^{-\delta}\widehat{\lambda}\right)^{\theta}}\right]^{1/\theta} \tag{46}$$

denotes the ratio of the effective intermediate input quality between the South and the North.

Proof. See Appendix B.7. □

The first three terms of (45) consist of parameters: agricultural fundamentals, which now include the land endowment, and non-agricultural productivity, which is now adjusted by the labor wedges. Over time, agricultural productivity differences therefore change only through the last two terms. The first is the South-to-North ratio of effective intermediate input quality, $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$, which enters with elasticity α, the production share of intermediate inputs. The second is the corresponding ratio of agricultural employment, which enters with elasticity $-\gamma$, where γ is the production share of land. This negative elasticity reflects congestion in agriculture: as more workers farm a fixed amount of land, agricultural labor productivity declines. I refer to the first as the effective input quality channel and the second as the structural change channel.

Equation (45) retains the effective input quality channel of the baseline expression (22) and adds the structural change channel. In the baseline model, relative agricultural employment is a function of the effective quality ratio alone, so substituting it into (45), with $\mathcal{Q}$ in place of $\widetilde{\mathcal{Q}}$ and without land endowments and labor wedges, gives the closed-form expression (22). In the quantitative model, trade costs make relative agricultural employment depend also on the two countries' expenditures on agricultural goods, which change through structural change, so agricultural productivity differences can no longer be expressed in closed form. I discuss the two channels in turn.

Effective quality of intermediate inputs. The ratio $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$ compares the effective quality of the intermediate inputs used by Southern farmers with that of the inputs used by Northern farmers. It generalizes the baseline ratio $\mathcal{Q}(\widehat{\lambda}_t)$ by capturing that farmers within a country use different input types. The baseline model features a bang-bang input choice: all farmers in a country use the same input type, so $\mathcal{Q}(\widehat{\lambda}_t) = \max\{\widehat{\lambda}_t, e^{-\delta}\}$. In the quantitative model, the ratio reflects the adoption shares in both countries, which vary smoothly with the price-adjusted quality ratio $\widehat{\lambda}_t$. As $\widehat{\lambda}_t$ declines, fewer farmers in each country use type-S inputs, and more use type-N inputs. Because type-N inputs are appropriate for the North but inappropriate for the South, the effective quality of the inputs used in the South then falls relative to that in the

North, and $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$ declines. As $\theta \to \infty$, the dispersion of suitability vanishes, all farmers in a country again use the same input type, and $\widetilde{\mathcal{Q}}$ converges to $\mathcal{Q}$.

The ratio $\widetilde{\mathcal{Q}}$ has two properties. First, it is a smooth and strictly increasing function of the price-adjusted quality ratio, with elasticity

$$\frac{\mathrm{d}\ln\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)}{\mathrm{d}\ln\widehat{\lambda}_t} = \omega_{S,S,t} - \omega_{N,S,t} > 0,$$

which is the difference between the adoption shares of type-S inputs in the South and the North. In the baseline model, by contrast, the elasticity of $\mathcal{Q}$ is zero for $\widehat{\lambda}_t < e^{-\delta}$, where no Southern farmer uses type-S inputs, and one for $\widehat{\lambda}_t > e^{-\delta}$, where all Southern farmers use them. For a given level of relative agricultural employment, a decline in the price-adjusted quality ratio therefore widens agricultural productivity differences at every level of this ratio. The effect is larger when the difference between the two adoption shares is larger. Second, as $\widehat{\lambda}_t$ approaches zero, the share of farmers using type-N inputs tends to one in both countries, and $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$ approaches its lower bound $e^{-\delta}$, the baseline value of $\mathcal{Q}$ when Southern farmers use inappropriate type-N inputs.

Relative agricultural employment. Relative agricultural employment changes over time through structural change. With a subsistence requirement, expenditure on agricultural goods includes the cost of subsistence consumption, $P_{j,t}^{a}\,\overline{a}\,L_j$ in (39). As the quality of intermediate inputs grows, agricultural goods become cheaper, and this cost falls relative to income. The expenditure share of agricultural goods therefore declines, and labor moves out of agriculture. This structural change proceeds at different speeds in the two countries. At the calibrated parameters, it is faster in the South, where the cost of subsistence consumption accounts for a larger share of expenditure on agricultural goods, so that a fall in their price lowers this expenditure proportionally more. Agricultural employment in the South then falls relative to the North, which raises Southern agricultural labor productivity relative to the North's.

Structural change affects relative agricultural employment only in the presence of trade costs. Without them, households in both countries face the same prices and therefore divide their expenditure between the two agricultural goods in the same proportions. Structural change then lowers agricultural employment by the same proportion in the two countries, regardless of which country's expenditure declines faster.[23] With trade costs, imported agricultural goods cost more, so households in each country spend relatively more on the domestic good. The demand for a country's agricultural good then depends more on that country's own expenditure on agricultural goods, and a faster decline of this expenditure in one country lowers its agricultural employment relative to the other's. In the limiting case without trade, this

[23] In this case, the demand for each country's agricultural good depends on world expenditure on agricultural goods, regardless of which country spends it. The division of world agricultural revenue between the two countries then depends only on the relative price of the two goods, which a decline in world expenditure leaves unchanged. Agricultural revenue and employment therefore fall by the same proportion in the two countries.

dependence is complete, because each country's agricultural revenue equals its own expenditure on agricultural goods.

Interaction between the two channels. The two channels are not independent, because relative agricultural employment also responds to the effective quality ratio. A higher effective quality ratio makes the Southern agricultural good cheaper, which has two effects that work in opposite directions. First, households substitute toward the Southern agricultural good, which raises Southern agricultural revenue and employment relative to the North's. Second, agricultural goods become cheaper in the South, which lowers the cost of subsistence consumption there and, for a given income, Southern expenditure on agricultural goods. This effect accelerates structural change in the South and lowers Southern agricultural employment relative to the North's. Without trade costs, only the first effect operates, because agricultural goods become cheaper in both countries alike. Without trade, only the second effect operates, because households cannot substitute toward the other country's agricultural good. At the calibrated parameters, the second effect dominates.

Land and the labor wedge. The land endowment is another component of agricultural fundamentals, together with exogenous agricultural productivity. The labor wedge lowers the agricultural wage, $w^a_{j,t} = A^n_j/(1+\tau_{w,j})$, which affects agricultural labor productivity in two ways. First, farmers' expenditure on intermediate inputs per worker is proportional to the agricultural wage, so a lower agricultural wage means fewer intermediate inputs per worker, thereby lowering agricultural labor productivity. This is why the ratio of agricultural wages, the third term of (45), enters with the exponent α. Second, farmers' labor demand implies $L^a_{j,t} = \beta(1+\tau_{w,j})\, p^a_{j,t} Y^a_{j,t}/A^n_j$. For a given level of agricultural revenue, a lower agricultural wage raises agricultural employment, which increases congestion in agriculture and lowers agricultural labor productivity. This effect operates through relative agricultural employment, as does the labor pull channel of Section 2.3.1.

4.3.2 Dynamic Lock-in through Directed Innovation

Dynamic lock-in. The mechanism of dynamic lock-in through directed innovation in Section 2.3.2 carries over to the quantitative model, with continuous adoption shares and innovation intensity replacing the bang-bang input choice and binary innovation decision. In the baseline model, lock-in unfolds through discrete changes, as type-S innovation stops altogether and all Southern farmers eventually switch to type-N inputs at once. In the quantitative model, innovation and adoption can decline gradually. As type-S inputs lag further behind the frontier, fewer farmers adopt them, reducing the market size for each improved product line. A lower quality ratio also provides greater scope for technological catch-up and can therefore support a higher margin rate. When the decline in market size dominates, operating profit $\pi_{S,t}(\lambda_{t-1})$ covers the R&D fixed costs of fewer product lines, and innovation intensity $m_{S,t}$

falls. If $m_{S,t} < \overline{m}(\lambda_{t-1})$, the quality ratio declines according to (42), reducing adoption of type-S inputs in period $t+1$ and potentially lowering innovation intensity further. Type-S inputs can thus fall further behind the frontier while some product lines are improved and many Southern farmers continue to use them. At sufficiently low quality ratios, operating profit falls below the lowest R&D fixed cost, and type-S innovation ceases permanently, as in the baseline model.

Other features of the quantitative model reinforce dynamic lock-in. Weak enforcement of intellectual property rights in the South and the additional cost of improving type-S inputs lower the profitability of type-S innovation at every quality ratio. Innovation intensity is therefore more likely to fall short of $\overline{m}(\lambda_{t-1})$, and type-$S$ inputs fall behind the frontier faster. In addition, structural change reduces agricultural revenue over time, so the market size for each improved product line shrinks even at a given quality ratio.

Dynamic amplification effect. As in Section 2.3.2, I measure how much directed innovation amplifies agricultural productivity differences by comparing the laissez-faire and counterfactual undirected innovation paths. On the latter, every product line of type-S inputs is improved in every period ($m_{S,t} = 1$), while all other equilibrium conditions remain in place. Both paths start from the same initial quality ratio $\lambda_{-1} \in (0, \chi]$. The first three terms of (45) are common to both paths, so the dynamic amplification factor is

$$\mathcal{A}_t := \frac{\text{APD}_t^U}{\text{APD}_t^D} = \left[\frac{\widetilde{\mathcal{Q}}(\widehat{\lambda}_t^U)}{\widetilde{\mathcal{Q}}(\widehat{\lambda}_t^D)}\right]^{\alpha} \left(\frac{L_{S,t}^{a,U}/L_{N,t}^{a,U}}{L_{S,t}^{a,D}/L_{N,t}^{a,D}}\right)^{-\gamma}, \tag{47}$$

where the superscripts U and D denote the undirected and laissez-faire paths, respectively. The first factor captures the contribution of the effective input quality channel, and the second captures that of the structural change channel. The structural change channel is new. In the baseline model, the amplification factor (23) depends only on the effective quality ratio. In the quantitative model, Southern farmers use more appropriate inputs on the undirected innovation path, which makes agricultural goods cheaper in the South and accelerates structural change there relative to the laissez-faire path. Fewer workers then remain in Southern agriculture, which reduces congestion in agriculture and raises Southern agricultural labor productivity. At the calibrated parameters, the second factor therefore exceeds one, so slower structural change in the South under laissez-faire adds to the amplification. Section 6 quantifies both factors and shows that structural change also shapes how agricultural productivity differences evolve along each path and how the subsidies affect welfare.

5 Calibration

This section calibrates the quantitative model of Section 4 to the US and Tanzania, representing the North and the South, respectively. Tanzania is a large agricultural economy in sub-Saharan

Africa. It is among the region's most populous countries, and agriculture employed about two-thirds of its workers in 2017.

One period is five years, and the calibration primarily targets 2017 data. I write λ_y for the quality ratio inherited at the start of the period beginning in year y. The initial ratio λ_{1982} corresponds to λ_{-1} in Section 4, and the inherited 2017 ratio is reached after seven periods of model dynamics. The dynamic amplification exercise in Section 6 compares two paths starting from the same quality ratio in 1982. Section 5.1 describes the external and internal calibration and the measurement of the inherited quality ratios. Section 5.2 evaluates the calibrated model's fit. Appendix C provides full calibration details.

5.1 Calibration Strategy

Table 1 summarizes the calibration. Panel A reports externally calibrated parameters, while Panel B reports parameters calibrated internally by matching model moments to the data, along with the quality ratio inherited in 2017. I measure this ratio from the varieties in use (Section 5.1.3) and take it as given for the parameters calibrated after it.

5.1.1 Externally Calibrated Parameters

Standard parameters. Following Restuccia et al. (2008) and Donovan (2021), I set the factor shares to $(\alpha, \beta, \gamma) = (0.40, 0.42, 0.18)$, consistent with US factor shares in gross agricultural output. I use US factor shares as a common benchmark for production elasticities in both countries. Holding these elasticities fixed across countries keeps the model tractable and focuses the analysis on the appropriateness of technology and directed innovation. I set the suitability dispersion parameter and the elasticity of substitution between the two agricultural varieties to $\theta = 2.46$ and $\varepsilon = 5.4$, respectively, using estimates from Costinot et al. (2016). The former is their estimated Fréchet shape parameter for productivity across fields, and the latter is their estimated elasticity of substitution between varieties of the same crop from different origins.

I calibrate the frontier quality step, $\xi = 0.177$, to US agricultural total factor productivity (TFP) growth over 1982–2017, using data from the US Department of Agriculture's Economic Research Service (ERS). Annual TFP growth of 1.30 percent implies a five-year growth factor of 1.067. In the model, frontier TFP grows entirely through input quality improvements and is proportional to quality raised to α, so $(1+\xi)^{\alpha} = 1.067$. I set the per-period discount factor to $1/(1+\rho) = 0.951$, corresponding to an annual discount rate of one percent, and utility curvature to $\sigma = 1$. These two parameters affect only the welfare calculations in Section 6.

I set relative non-agricultural productivity to $A_S^n / A_N^n = 0.167$, using the 2017 ratio of non-agricultural value added per person engaged, measured at sectoral purchasing power parities (PPPs), from the GGDC Productivity Level Database (PLD) (Inklaar et al., 2023). The relative land endowment is $Z_S / Z_N = 0.093$, based on the 2017 ratio of cropland (arable land and land under permanent crops) from FAOSTAT. I do not adjust cropland for quality, so land quality

Table 1: Model calibration

Parameter	Description	Value	Source
Panel A. Externally calibrated parameters			
α, β, γ	Factor shares (inputs, labor, land)	0.40, 0.42, 0.18	Restuccia et al. (2008), Donovan (2021)
θ	Fréchet parameter	2.46	Costinot, Donaldson, and Smith (2016)
ε	Agricultural Armington elasticity	5.4	Costinot et al. (2016)
ξ	Frontier quality step size	0.177	US agricultural TFP growth (1982–2017)
ρ, σ	Discount rate, Utility curvature	0.0515, 1	1% per year, Standard
δ	Environmental mismatch penalty	1.45	Moscona and Sastry (2025)
A_S^n / A_N^n	Non-agricultural productivity ratio	0.167	Non-agricultural VA per worker (2017)
Z_S / Z_N	Land endowment ratio	0.093	Cropland (2017)
L_S / L_N	Labor endowment ratio	0.130	Persons engaged (2017)
ζ_S	Enforcement of IP rights	0.78	Campi and Nuvolari (2021)
$1+\tau_{w,N}, 1+\tau_{w,S}$	Labor wedges	1.75, 21.9	Sectoral gaps in VA per worker (2017)
τ_g	Trade cost on agricultural goods	4.58	Own-supply share of food (2017)
f	Additional type-S R&D fixed cost	0.34	Phillips McDougall (2011, 2016); AgbioInvestor (2024, 2026)

Panel B. Internally calibrated parameters

Parameter	Description	Value	Target moment	Data	Model
ϕ	Agricultural expenditure weight	0.0137	US agricultural employment share (2017)	0.0137	0.0137
η	Research productivity	4338	US private agricultural-input R&D intensity (2010–2014)	0.0124	0.0124
A_S^a	Exogenous agricultural productivity ($A_N^a = 1$)	0.118	Agricultural productivity differences (2017)	1/160	1/160
$\bar{a}$	Subsistence consumption per capita	0.0165	East Africa's share of agricultural output (2017)	0.156	0.156
b	R&D fixed-cost distribution half-width	0.73	Private breeder presence in East Africa (TASAI, 2021)		
a, χ	Adaptation rate, Transferability	0.29, 0.65	Modern variety adoption during the Green Revolution in Evenson and Gollin (2003)		
λ_{2017}	Quality ratio inherited in 2017	0.236	Modern variety shares and ages of varieties in use (Section 5.1.3)		

Notes: Panel A lists the externally calibrated parameters with their sources, and Panel B the internally calibrated parameters with their target moments (the data value and the model value) and the quality ratio inherited in 2017. The North is the US, and the South is Tanzania.

differences are absorbed by the agricultural productivity parameter A_S^a. The relative labor endowment is $L_S/L_N = 0.130$, based on the 2017 ratio of persons engaged from the PLD.

The environmental mismatch penalty. I calibrate δ in (2) using the micro estimates of Moscona and Sastry (2025). They construct a crop-level index of pest and pathogen dissimilarity between countries, ranging from zero (no mismatch) to one (maximum mismatch), and estimate its effect on agricultural output. First, I compute Tanzania's harvested-area-weighted mean mismatch with the US across crops, obtaining 0.367. Second, their estimated log-output loss of 9.71 per unit of mismatch includes crop-choice responses. Dividing by their calibrated Fréchet shape parameter, 2.46, removes these responses and yields a log productivity loss of 3.95 per unit of mismatch. Third, multiplying this loss by the mismatch index gives $\delta = 1.45$, so frontier inputs retain $e^{-\delta} = 0.234$ of their quality in Tanzania before adaptation. This mapping interprets their crop-level productivity loss as a loss of effective input quality from environmental mismatch.

I follow Moscona and Sastry (2025) in treating the US as the technological frontier, given its dominance in agricultural science and patenting. Their baseline defines the frontier crop by crop as the countries releasing the most varieties, but their estimates are similar when the US serves as the frontier for every crop. I use the estimate from the latter specification. The penalty is symmetric in (2), so type-S inputs face the same proportional quality loss in the North. This matters little because Northern farmers mainly use type-N inputs.

Enforcement of intellectual property rights. The parameter ζ_S in (32) is the share of operating profits from sales to Southern farmers retained by the patent holder. I set it equal to the 2018 index of intellectual property protection for plant varieties from Campi and Nuvolari (2021), normalized by its maximum, so that values closer to one indicate stronger protection. The index scores five features of a country's statutes: the acts of the International Union for the Protection of New Varieties of Plants that it has ratified, the farmers' exception, the breeders' exception, the duration of protection, and the scope of patent protection in agriculture. This gives $\zeta_S = 0.78$.

The additional cost of adapting inputs to the South. The parameter f in (34) captures the additional R&D fixed cost of improving type-S product lines relative to type-N product lines. I interpret it as the share of the cost of developing a frontier input that must be incurred again to adapt it to Southern conditions. I construct it from published cost breakdowns for new agricultural inputs (Phillips McDougall, 2011, 2016; AgbioInvestor, 2024, 2026), which distinguish research, development, and registration. I count field trials and environment-specific development studies as costs that must be incurred again for adaptation.[24] The implied values of f

[24] I treat registration costs for conventional varieties as negligible. The African Seed Access Index (TASAI) reports release costs of $339 to $41,668 per variety and jurisdiction in East Africa (Mabaya et al., 2020, 2025; Mburu et al., 2022; Mabaya, Mizambwa, et al., 2021; Mabaya et al., 2022; Mabaya, Hassena, et al., 2021; Mabaya, Urinzwenimana, et al., 2021). These are three to five orders of magnitude below the $33 to $42 million in registration costs

range from 0.16 to 0.51, and I use the midpoint, $f = 0.34$. Alternative calculations that retain registration costs for regulated traits or include all development routes yield midpoints of 0.41 and 0.39, respectively. Appendix C.2 describes the construction in detail.

The labor wedge and the trade cost. I calibrate these two parameters directly from observed moments using model identities. For the labor wedge $\tau_{w,j}$ in (41), I use the within-country ratio of non-agricultural to agricultural value added per worker, which is the agricultural productivity gap studied by Gollin, Lagakos, and Waugh (2014a) and Herrendorf and Schoellman (2018). In the model, agricultural labor demand is $L^a_{j,t} = \beta(1+\tau_{w,j})\, p^a_{j,t} Y^a_{j,t} / A^n_j$, while agricultural value added is $(1-\alpha)\, p^a_{j,t} Y^a_{j,t}$. These imply

$$\frac{\beta(1+\tau_{w,j})}{1-\alpha} = \frac{\text{non-agricultural value added per worker in country } j}{\text{agricultural value added per worker in country } j}. \tag{48}$$

I construct the right-hand side using 2017 data consistent with the model's definition of value added. For non-agriculture, I use value added per person engaged from the PLD. For agriculture, I multiply FAO production value at current producer prices by $(1-\alpha)$ and divide by agricultural employment from the PLD. The US gap of 1.23 gives $1+\tau_{w,N} = 1.75$, while Tanzania's gap of 15.3 gives $1+\tau_{w,S} = 21.9$.

I calibrate the trade cost τ_g to the share of Southern agricultural expenditure allocated to its own variety. With equal weights, the price index in (38) implies

$$\frac{p^a_{S,t}\, c^a_{S,S,t}}{E^a_{S,t}} = \frac{1}{1 + \left(p^a_{S,t} / (\tau_g\, p^a_{N,t})\right)^{\varepsilon-1}}, \tag{49}$$

which I invert for τ_g using the farm-gate price ratio $p^a_{S,t}/p^a_{N,t}$ and the observed own-supply share. I measure the price ratio as the PLD's agricultural output parity, based on farm-gate producer prices, divided by its non-agricultural parity. For the own-supply share, I use one minus the ratio of food imports from high-income countries to apparent consumption (production plus imports minus exports), using UN Comtrade and FAO data for 2017. The price ratio of 2.14 and the own-supply share of 0.966 yield $\tau_g = 4.58$.

5.1.2 Internally Calibrated Parameters

The internal calibration is recursive. First, I calibrate the catch-up parameters a and χ, which govern the evolution of the quality ratio, to historical adoption data. Second, given these parameters, I measure the quality ratio inherited in 2017 from the varieties in farmers' fields (Section 5.1.3). Third, conditional on this ratio, I calibrate four parameters to four moments. Fourth, I set the dispersion of R&D fixed costs across product lines using evidence on whether private

reported for frontier crop-protection products.

breeders are active in East African seed markets. For ease of exposition, I first present the measurement of agricultural productivity differences and the calibration to the four moments. I then turn to the catch-up law and the dispersion of R&D fixed costs.

Measuring agricultural productivity differences. The central target is APD_{2017}, the 2017 ratio of agricultural output per worker in Tanzania to that in the US. Since the production function (27) specifies gross output, I use FAO's gross production value of agriculture at constant 2014–2016 international prices. These prices value both countries' crop and livestock output on a common basis, so output differences reflect quantities. This follows the output concept used in the ERS's International Agricultural Productivity accounts (USDA Economic Research Service, 2026) and by Restuccia et al. (2008). I measure agricultural employment using persons engaged from the PLD, maintaining the same workforce concept as for A_S^n/A_N^n. This gives $\text{APD}_{2017} = 1/160$.[25,26]

Four parameters from four moments. Given the externally calibrated parameters, the catch-up parameters, and the quality ratio inherited in 2017, I calibrate four remaining parameters to four moments. First, I choose the agricultural expenditure weight ϕ to match the US agricultural employment share of 1.37 percent in 2017. A higher ϕ raises Northern agricultural revenue and farmers' labor demand. Matching this share gives $\phi = 0.0137$. Second, I calibrate research productivity η to US private agricultural-input R&D intensity: 1.24 percent of agricultural production value, based on ERS data for 2010–2014. In the model, this ratio is $c_N/(p_{N,t}^a Y_{N,t}^a)$, where $c_N = w_{N,t}^r/\eta$. The ratio falls with η, and matching the target gives $\eta = 4338$. Third, I calibrate exogenous agricultural productivity in the South, A_S^a, with $A_N^a = 1$, to match APD_{2017}. Inverting Proposition 6 at the inherited 2017 quality ratio and the period's equilibrium innovation intensity gives $A_S^a = 0.118$. Fourth, I calibrate subsistence consumption $\bar{a}$ to East Africa's share of combined regional and US agricultural output, measured at constant international prices. I aggregate Ethiopia, Kenya, Malawi, Rwanda, Tanzania, and Uganda into one Southern economy.[27] Matching its 2017 output share of 0.156 yields $\bar{a} = 0.0165$.

[25] I construct this ratio, the non-agricultural productivity ratio, the within-country gaps in (48), and the farm-gate price ratio in (49) from consistent production values, workforce counts, and price indices. The model's accounting identity linking these four objects therefore holds exactly in the data. A common workforce concept matters because its choice affects the productivity ratio more than the choice of price basis.

[26] Agricultural productivity differences in gross output per worker are larger than those in value added per worker in the PLD. In Tanzania, national-accounts agricultural value added exceeds the farm-gate value of crop and livestock production, whereas US agricultural value added is about half of gross output after deducting purchased inputs. Appendix C.1 explains the relationship between the two measures.

[27] I use the regional target to reduce sensitivity to measurement error in small country-level output shares. This also leaves Tanzania's agricultural employment share and its share of the pair's agricultural revenue untargeted, as reported in the first two rows of Table 2. Subsistence consumption is needed to match the data. Under homothetic preferences ($\bar{a} = 0$), the South's small share of world income limits agricultural revenue. The model then predicts a Southern agricultural employment share far below the data and a regional output share of 0.022, compared with

The catch-up law. The adaptation rate a and transferability ceiling χ in (33) determine an improved type-S product line's quality relative to the frontier: $T(\lambda_{t-1}) = (1-a)\lambda_{t-1} + a\chi$. If every product line is improved in every period, the quality ratio converges to χ, closing a fraction a of the remaining distance each period. Closing 95 percent of the initial distance therefore takes $5\ln(0.05)/\ln(1-a)$ years.

I calibrate a and χ using the Green Revolution in India, Pakistan, Bangladesh, Indonesia, the Philippines, and Mexico, where breeding programs improved major food-crop varieties from the mid-1960s onward. I represent these episodes as the model's case of innovation in every type-S product line in every period. Although quality ratios are unobserved, modern-variety adoption provides an observable counterpart. Evenson and Gollin (2003) report harvested-area shares under varieties released after 1965, by region and crop, for 1970, 1980, 1990, and 1998. Weighting these crop shares by each economy's harvested areas gives four adoption observations per economy. For each country j, I map modern varieties to type-S inputs with relative quality $\lambda_{j,t}$ and assign traditional varieties a relative quality equal to the mismatch penalty $e^{-\delta_j}$, the lowest value consistent with farmers' choices. Equation (30) then gives

$$\omega_{j,t} = \frac{(\lambda_{j,t}e^{\delta_j})^\theta}{1 + (\lambda_{j,t}e^{\delta_j})^\theta}.$$

I invert this equation to obtain the quality ratio $\lambda_{j,t}$ implied by the observed adoption share $\omega_{j,t}$.

I fit the catch-up paths from 1965 to these adoption shares, allowing an economy-specific initial quality ratio and measuring residuals in log points of the quality ratio. I also allow adoption to adjust gradually toward the model-implied share, fitting a separate adjustment lag for each economy. This prevents delayed adoption from being interpreted as slow technological improvement. The fit does not sharply distinguish between parameter pairs. For each ceiling χ, there is a best-fitting speed $a^*(\chi)$, and the fit varies little along this curve. Two bounds restrict its admissible range. First, adoption at any date cannot exceed the share implied by $\lambda_{j,t} = \chi$, so India's 1998 share of 0.72 requires $\chi \geq 0.63$. Second, I restrict the time to close 95 percent of the distance to the ceiling to 25–50 years, based on historical catch-up episodes. This bounds a between 0.26 and 0.45. Because $a^*(\chi)$ declines with χ, it also places an upper bound of 0.72 on χ along the fitted curve. I choose $\chi = 0.65$ and its best-fitting speed, $a = 0.29$. At the chosen pair, closing 95 percent of the distance to the ceiling takes 44 years, with a half-life of 10 years. Appendix C.10 shows that the dynamic amplification factor changes little near both ends of the admissible range.

The dispersion of R&D fixed costs. The fixed cost of improving a type-k product line, $c_{k,\ell,t}$ in Section 4.1.2, is uniformly distributed over $[(1-b)c_k, (1+b)c_k]$. The lowest cost for type-S inputs is therefore $(1-b)c_S$, so private innovation in type-S inputs starts once the profit-to-cost ratio $\pi_{S,t}/c_S$ reaches $1-b$. I calibrate b from the presence of private breeders, which reveals on

0.156 in the data.

which side of this threshold an economy stands: an economy with private breeders has a profit-to-cost ratio at or above $1 - b$, and an economy without them has a ratio below it. The TASAI country reports count seven private breeders among 36 active breeders in Tanzania, and none in neighboring Rwanda, whose seed sector the same reports describe (Mabaya, Mizambwa, et al., 2021; Mabaya, Urinzwenimana, et al., 2021). The threshold $1 - b$ therefore lies between the two economies' profit-to-cost ratios, which the calibrated model puts at 0.218 in Rwanda and 0.323 in Tanzania in 2017. I set it at the geometric mean of the two ratios, 0.27, which gives $b = 0.73$. This choice is the midpoint in proportional terms, so a small error in either calibrated ratio does not move an observation across the threshold. The calibration yields a 2017 type-S innovation intensity of 0.039, meaning that 3.9 percent of type-S product lines are improved. This value lies below $\overline{m}(\lambda_{2017}) = 0.284$, the intensity needed to hold the quality ratio constant in (43). The quality ratio therefore falls under laissez-faire. These conditions provide the starting point for the quantitative analysis in Section 6.

5.1.3 The Quality Ratio in 2017 and in 1982

The quality ratio in 2017. Since variety quality relative to the frontier is not directly observed, I infer it from two crop-level measures: the area share under modern varieties, from the Diffusion and Impact of Improved Varieties in Africa (DIIVA) database, with fieldwork conducted in 2009–2011, and the age of varieties on sale in 2020, from the TASAI country report. A Southern variety released when the quality ratio was λ_r has quality $T(\lambda_r)$ relative to the frontier at release. I adjust this quality for the age gap relative to varieties used by US farmers, dividing it by $1 + \xi$ for every five years of this gap. I assign landraces the relative quality $e^{-\delta}$, the lowest value consistent with farmers' choices. Following the knowledge pooling in (31), I average log qualities using harvested-area weights. Setting the unobserved release-date ratio λ_r to λ_{2017} itself yields an equation with a unique solution between $e^{-\delta}$ and χ.

This gives $\lambda_{2017} = 0.236$, only 0.7 percent above $e^{-\delta}$. This low value reflects the age of Tanzania's varieties: maize varieties on sale are 22.8 years old on average, compared to 4 years for those used in the US. After adjusting for age, relative quality is no higher than the mismatch penalty $e^{-\delta}$ on 71 percent of the area under modern varieties. This measurement uses no yield, price, or output data, leaving agricultural productivity differences to be matched by A_S^a. Appendix C.4 describes the construction in detail.

The quality ratio in 1982. The 1982 quality ratio is unobserved. I infer it as the initial state from which the laissez-faire equilibrium path reaches the measured λ_{2017} after seven five-year periods. All parameters remain at their calibrated values. Since subsistence consumption is fixed in units of agricultural goods, the equilibrium of each period also depends on the level of frontier quality, which is lower in earlier periods. The law of motion (42) gives

$$\lambda_{y+5} = \lambda_y \exp\{m_{S,y}(\lambda_y) \ln G_S(\lambda_y) - \ln(1 + \xi)\},$$

where $m_{S,y}(\lambda_y)$ is the equilibrium type-S innovation intensity in the period beginning in year y, given the inherited quality ratio λ_y, and G_S is the gross step size in (18).

I recover λ_{1982} by working backward one period at a time. Given λ_{y+5}, I solve this equation for λ_y on the interval $[\lambda_{y+5}, (1+\xi)\lambda_{y+5}]$. This interval reflects that the quality ratio does not rise along this path and falls most rapidly when no type-S innovation occurs. If no type-S product line is improved throughout 1982–2017, the solution simplifies to $\lambda_{1982} = (1+\xi)^7 \lambda_{2017} = 3.12\,\lambda_{2017}$. When some product lines are improved, the quality ratio declines more slowly, implying a lower initial ratio than the no-innovation formula. Along this path, type-S innovation occurs in every period, with intensity no higher than 0.083, giving $\lambda_{1982} = 0.631$.

5.2 Model Fit

The calibration targets four moments and Green Revolution adoption curves. Table 2 compares the calibrated model with untargeted moments. The first three rows concern 2017. The Southern agricultural employment share is untargeted because subsistence consumption $\bar{a}$ is calibrated to a regional output share. The model predicts a share of 0.600, compared with 0.691 in the data. The South's share of the pair's agricultural revenue is untargeted for the same reason. Its data counterpart converts the quantity share at international prices into a value share with the farm-gate price ratio, and the model's 0.071 is close to the 0.081 in the data. GDP per worker relative to the US compares the model's income per worker on the GDP concept with the PLD's GDP per person engaged at the non-agricultural PPP. It is the object of the exercise of Restuccia et al. (2008) and is reproduced without being targeted, the model's 0.073 compared with 0.073 in the data.

The remaining rows assess the untargeted dynamics since 1982. The data are the ERS's International Agricultural Productivity accounts, which value gross agricultural output at constant international prices, as the measurement of Section 5.1.2 does, and count the adults whose main economic activity is agriculture, a workforce concept that differs from the PLD's persons engaged. The two concepts give productivity ratios within 1.2 percent of each other in 2015–2019. In these accounts, agricultural output per worker in Tanzania relative to the US averaged 1/123 in 1980–1984, compared with 1/120 along the model's laissez-faire path in 1982.[28] Over 1982–2017, the observed ratio falls to 0.78 of its 1982 value, the mean of 2015–2019 over the mean of 1980–1984. The corresponding value under laissez-faire is 0.75. Under undirected innovation, in which both input types improve every period, the ratio instead rises to 1.06 times its 1982 value. The laissez-faire path thus captures both the direction and magnitude of the observed change.

[28] Restuccia et al. (2008), whose 1985 benchmark falls within the model's first period, give 1/119 for Tanzania. They value output at the international prices of their benchmark year and measure the workforce by FAO's agricultural employment series, so that comparison concerns the order of magnitude rather than the level.

Table 2: Model fit to untargeted moments

Moment	Model	Data
Southern agricultural employment share in 2017	0.600	0.691
Southern share of the pair's agricultural revenue in 2017	0.071	0.081
GDP per worker relative to the US in 2017	0.073	0.073
APD_t in 1982	1/120	1/123
APD_t in 2017 relative to 1982	0.75	0.78
the same on the undirected path	1.06	–

Notes: The first three rows are moments of 2017: the Southern agricultural employment share, the South's share of the agricultural revenue of the pair, and GDP per worker relative to the US at the non-agricultural PPP. The remaining rows concern the dynamics: agricultural productivity differences (agricultural output per worker in Tanzania relative to the US) in the first period of the model compared to the mean of 1980–1984 in the ERS's International Agricultural Productivity accounts, and the ratio of agricultural productivity differences in 2017 to those in 1982 in the model (the laissez-faire path and the undirected path) compared to the same ratio in those accounts (the mean of 2015–2019 over the mean of 1980–1984).

6 Quantitative Results

This section uses the calibrated model to answer the paper's two central questions: how much directed innovation contributes to agricultural productivity differences between Tanzania and the US, and how effective policy interventions are. Section 6.1 answers the first question by measuring the dynamic amplification effect. Sections 6.2 to 6.4 answer the second by evaluating three policies: an R&D subsidy that holds the quality ratio at its 2017 value (Section 6.2), a larger R&D subsidy under which every product line of type-S inputs is improved (Section 6.3), and stronger enforcement of intellectual property rights in the South, followed by further counterfactuals (Section 6.4). In each case, the economy follows the laissez-faire path from 1982 until the policy is introduced in 2017. I write t_0 for the period beginning in 2017 and APD_t for agricultural labor productivity differences between the South and the North, $(Y^a_{S,t}/L^a_{S,t})/(Y^a_{N,t}/L^a_{N,t})$.

6.1 The Dynamic Amplification Effect

I compare two paths of the calibrated economies from 1982. The first is the laissez-faire equilibrium path, on which R&D firms direct innovation toward the input type that is profitable to improve. The second is the counterfactual undirected innovation path of Section 2.3.2, on which every product line of type-S inputs is improved in every period ($m_{S,t} = 1$), as every product line of type-N inputs is on both paths. Both paths start from the same quality ratio in 1982, and all other parameters are the same. The dynamic amplification factor $\mathcal{A}_t$ in (47) is the ratio of agricultural productivity differences on the undirected path to those on the laissez-faire

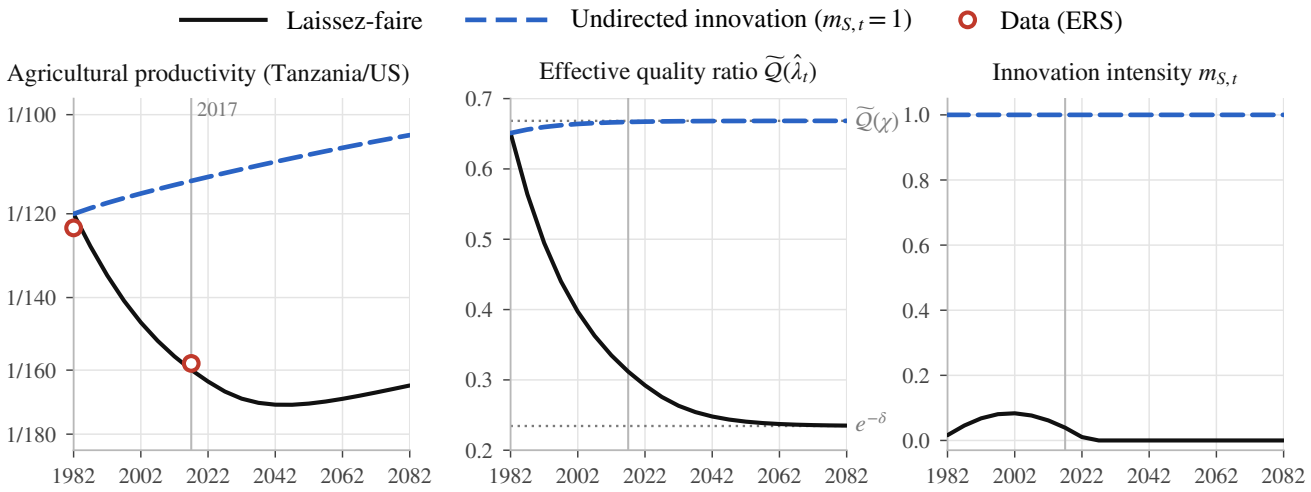


Figure 2: The dynamic amplification effect for Tanzania and the US: the laissez-faire path (solid lines) and the counterfactual undirected innovation path (dashed lines) from 1982. The left panel plots the Tanzania-to-US agricultural labor productivity ratio, APD_t. The middle panel plots the effective quality ratio, $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$, with its lower bound $e^{-\delta}$ and upper bound $\widetilde{\mathcal{Q}}(\chi)$ (dotted lines). The right panel plots the innovation intensity of type-S inputs, $m_{S,t}$. On the undirected path, every product line of type-S inputs is improved in every period ($m_{S,t} = 1$). The vertical line marks 2017. The circles show agricultural output per worker in Tanzania relative to the US in the ERS's International Agricultural Productivity accounts, averaged over 1980–1984 and 2015–2019; the 2017 ratio is the calibration's target, and the 1982 ratio is untargeted (Section 5.2).

path.

Figure 2 shows the two paths from 1982 to 2082. The left and middle panels plot the Tanzania-to-US ratios of agricultural labor productivity and effective input quality, respectively, where the effective quality ratio, $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$, is defined in equation (46). The right panel plots the innovation intensity of type-S inputs, $m_{S,t}$. The circles in the left panel show data from the ERS's International Agricultural Productivity accounts, averaged over 1980–1984 and 2015–2019. The 2017 agricultural productivity ratio corresponds to the calibration target, with 1/160 in the model compared with 1/158 in the ERS data. The 1.2 percent difference reflects the ERS accounts' use of a different agricultural workforce measure. The 1982 ratio is untargeted: the laissez-faire path gives 1/120, compared with 1/123 in the data.

Under laissez-faire, the effective quality ratio falls from 0.65 in 1982 to 0.31 in 2017 and 0.24 in 2047 (middle panel). This reflects a decline in the underlying quality ratio of type-S to type-N inputs, which equals 0.63, 0.24, and 0.09 in those years, respectively. Type-S innovation occurs, but at an intensity far below $\overline{m}(\lambda_{t-1})$ in equation (43): the intensity is 0.016 in 1982, reaches 0.08 in 2002, falls to 0.04 in 2017, and is zero from 2027 onward. This is the second case of Section 4.2: some product lines of type-S inputs are improved while type-S inputs fall further behind the frontier. Agricultural labor productivity in Tanzania relative to the US falls from 1/120 in 1982

to 1/160 in 2017, reaches its lowest value, 1/171, in 2047, and then rises slowly, to 1/165 in 2082.

Two channels shape the laissez-faire path. By (45), agricultural productivity differences change over time only through the effective quality ratio $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$, with the elasticity α, and through agricultural employment in the South relative to the North, with the elasticity $-\gamma$: the effective input quality channel and the structural change channel of Section 4.3.1. The effective input quality channel accounts for the fall. In 1982, type-S inputs account for a share 0.92 of Southern farmers' expenditure on intermediate inputs, and their quality falls further behind the frontier in every period. As the price-adjusted quality ratio declines, farmers switch to type-N inputs, which are adapted to Northern environments and carry the mismatch penalty: by 2047, the expenditure share of type-S inputs is 0.09. Between 1982 and 2047, this channel lowers APD_t by 32.5 percent. The switch also bounds the channel. The effective quality ratio $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$ in (46) cannot fall below $e^{-\delta} = 0.23$, the lower bound in the middle panel, so once Southern farmers use type-N inputs, effective input quality no longer widens agricultural productivity differences.

The structural change channel accounts for the subsequent rise. Subsistence consumption is fixed in units of agricultural goods, whereas agricultural output grows with the quality of intermediate inputs (Section 4.2). As the frontier quality rises, agricultural goods become cheaper, the cost of subsistence consumption falls relative to income, and workers leave agriculture. The farm employment share falls far more in the South, where the cost of subsistence consumption accounts for most of the expenditure on agricultural goods, than in the North: from 0.73 in 1982 to 0.45 in 2047 and 0.21 in 2132 in the South, and from 0.016 to 0.011 in the North. With land fixed, fewer workers on the same land reduce congestion in agriculture and raise output per farm worker, by more in the South than in the North, so agricultural productivity differences narrow.

The two channels are not independent. Until the switch to type-N inputs is complete, the effective quality of the inputs that Southern farmers use grows more slowly than the frontier quality, so agricultural goods become cheaper more slowly in the South than in the North: between 1982 and 2047, the price index of agricultural goods falls by 0.8 percent per year in the South, against 1.4 percent in the North. The cost of subsistence consumption therefore remains high in the South, and structural change is slow: over these years, the structural change channel raises APD_t by 4.3 percent, against 9.6 percent on the undirected path. Once Southern farmers use type-N inputs, the effective quality of their inputs grows at the same rate as the frontier quality, agricultural goods become cheaper at least as fast as in the North, and structural change accelerates: between 2047 and 2082, the channel raises APD_t by 5.1 percent.

The innovation intensity in the right panel is hump-shaped. It follows the operating profit of an improved product line of type-S inputs, which is the product of the margin rate and the market size (Section 4.3.2), and the two move in opposite directions as type-S inputs fall behind the frontier. In 1982, the quality ratio is close to the transferability ceiling, so there is little scope for technological catch-up, and the margin rate is low. As the quality ratio falls, the scope for

catch-up widens, and the margin rate rises, from 0.16 in 1982 to 0.32 in 2002 and 0.44 in 2017. Until 2002, this force dominates, and the intensity rises. The market size shrinks throughout, for two reasons. First, farmers switch to type-N inputs: the expenditure share of type-S inputs in the South falls from 0.92 in 1982 to 0.50 in 2017. Second, structural change contracts Southern agriculture: Southern agricultural revenue falls by 17 percent over the same period. After 2002, the decline in market size dominates, and the intensity falls. From 2027 onward, the operating profit is below the lowest R&D fixed cost, and no product line of type-S inputs is improved.

On the undirected path, the quality ratio rises from 0.63 toward the transferability ceiling $\chi = 0.65$ within a few periods, and the effective quality ratio from 0.65 to its upper bound $\widetilde{\mathcal{Q}}(\chi) = 0.67$ (middle panel). Southern farmers keep using type-S inputs that keep pace with the frontier, so the effective input quality channel is nearly inactive, and Tanzania's relative agricultural productivity rises from the start through the structural change channel: 1/113 in 2017, 1/108 in 2047, and 1/104 in 2082. The ratio of the undirected path to the laissez-faire path is the dynamic amplification factor. It equals 1.42 in 2017, 1.55 in 2037, and 1.57 in 2047. Directed innovation thus amplifies agricultural productivity differences between Tanzania and the US by 42 percent in 2017, rising to 57 percent by 2047.

Both channels contribute to the amplification, and (47) separates their contributions. In 2017, the effective quality ratio is 0.31 on the laissez-faire path against 0.67 on the undirected path, which accounts for a factor of 1.36. The remaining factor of 1.04 reflects slower structural change: because agricultural goods are more expensive in the South under laissez-faire, the farm employment share of the South in 2017 is 0.60, against 0.47 on the undirected path. The first part is permanent, and the second is not. In the long run, both paths converge to constant agricultural productivity differences, 1/140 under laissez-faire and 1/93 on the undirected path, a ratio of 1.51 (Appendix D.5 shows the paths to 2382). The labor allocations of the two paths become nearly the same once subsistence consumption no longer matters, whereas the quality ratio converges to zero on the laissez-faire path and to χ on the undirected path: the South under laissez-faire permanently uses type-N inputs with the mismatch penalty. Because the second part fades, the factor declines after reaching its highest value, 1.59, in 2067, toward its long-run value, which is close to $(\widetilde{\mathcal{Q}}(\chi)/e^{-\delta})^{\alpha} = 1.52$, the contribution of the effective input quality channel alone. Most of the amplification is thus a permanent level effect rather than a transition.

6.2 An R&D Subsidy That Holds the Quality Ratio

The subsidy. The government pays a subsidy $s_t \geq 0$ to the R&D firm of type-S inputs for every product line it improves in period t. The manager of product line ℓ then improves it if and only if $\pi_{S,t} + s_t \geq c_{S,\ell,t}$, so the innovation intensity of type-S inputs in (35) becomes

$$m_{S,t} = H\left(\frac{\pi_{S,t}(\lambda_{t-1}) + s_t}{c_S}\right),$$

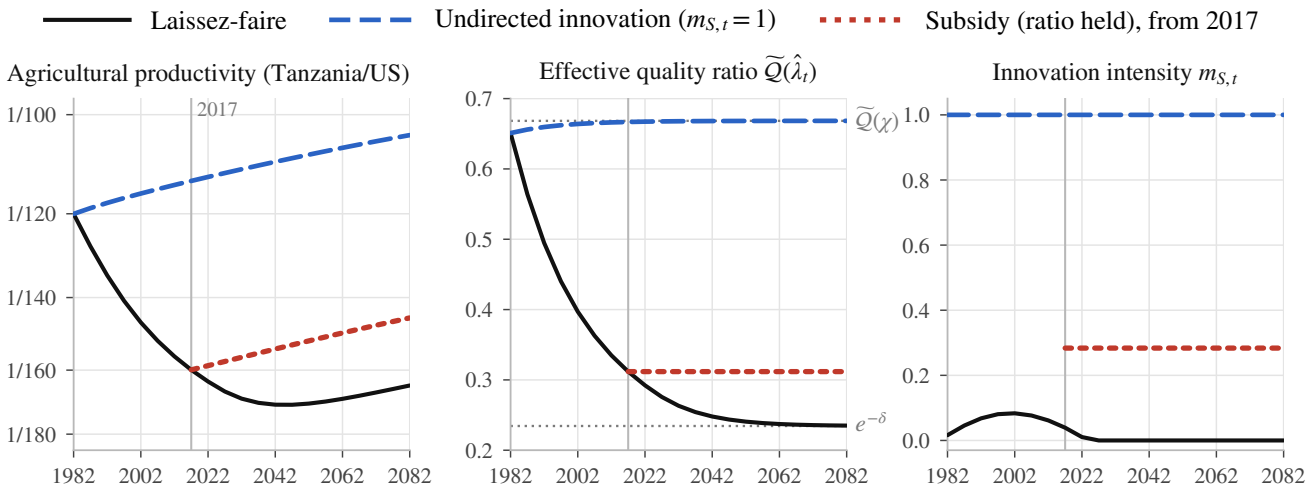


Figure 3: The R&D subsidy that holds the quality ratio at its 2017 value: the laissez-faire path and the undirected path from 1982, as in Figure 2, and the path under the subsidy from 2017, Tanzania and the US. The panels are those of Figure 2. Southern households finance the subsidy.

and the subsidy receipts $s_t m_{S,t}$ are part of the net profit of the R&D firm, which is located in the North. The subsidy is financed by a lump-sum tax $\tau_{S,t} = s_t m_{S,t}$ on Southern households, whose income is $E_{S,t} - \tau_{S,t}$, with $E_{S,t}$ given by (40). The South thus pays Northern R&D firms to improve the inputs adapted to its conditions. Appendix D.4 considers the alternative in which Northern households finance the subsidy.

The headline program holds the quality ratio at its 2017 value. In every period $t \geq t_0$, the government sets the subsidy rate at which the innovation intensity equals $\overline{m}(\lambda_{t_0-1})$, the intensity at which the quality ratio remains constant in (43), so that $\lambda_t = \lambda_{t_0-1}$ for all $t \geq t_0$.[29] At the calibrated parameters, the required intensity is 0.28, against a laissez-faire intensity of 0.04 in 2017 and zero from 2027 onward. The subsidy rate in 2017 equals 0.36 of the mean R&D fixed cost of a product line of type-S inputs, and the cost of the subsidy, its total payment $s_t m_{S,t}$, equals 0.33 percent of Southern GDP, which is household income before the tax, or 2.2 percent of Southern agricultural revenue.[30] The rate rises slowly over time, to 0.52 of the mean fixed cost by 2082. Because the intensity is held constant, the rise in the rate equals the fall in the operating profit of a product line of type-S inputs at the held quality ratio. At this quality ratio, the margin rate and the adoption shares do not change, so the operating profit falls only because structural change contracts agriculture: Southern agricultural revenue is 50 percent lower in 2082 than in 2017.

[29] If the laissez-faire intensity exceeded $\overline{m}(\lambda_{t_0-1})$ in some period, the subsidy rate would be zero in that period and the quality ratio would rise. This does not occur at the calibrated parameters: the laissez-faire intensity is below $\overline{m}(\lambda_{t_0-1})$ in every period, so the subsidy is paid in every period.

[30] The costs of this section are in percent of Southern income before the tax, evaluated on the equilibrium path under the subsidy, which Southern households finance. In percent of their income after the tax, the same costs are 0.33 and 4.75 percent.

Productivity. Figure 3 shows the laissez-faire path and the undirected path of Figure 2 from 1982, and the path under the subsidy from 2017 to 2082. The subsidy holds the quality ratio at 0.24, and with it the effective quality ratio at 0.31, while under laissez-faire both keep falling (middle panel); it holds the innovation intensity at 0.28 (right panel). The subsidy thus switches off the effective input quality channel. Type-S inputs keep pace with the frontier, their share in Southern farmers' expenditure on intermediate inputs stays at 0.50, and from 2017 onward only the structural change channel operates, so agricultural productivity differences narrow (left panel). Agricultural labor productivity in Tanzania relative to the US rises above the laissez-faire path from 2022 onward, one period after the subsidy starts: 1/155 against 1/170 in 2037 and 1/153 against 1/171 in 2047, a rise of 9.6 percent in 2037 and 11.6 percent in 2047. Of the 11.6 percent in 2047, 10.4 percentage points reflect the higher effective quality ratio under the subsidy, and the remainder reflects faster structural change. The rise grows for several decades, to 13.2 percent in 2082, and changes little thereafter. It grows while the laissez-faire South is still switching to type-N inputs and its effective quality ratio is still falling. Once the switch is complete, the effective quality ratio is constant on both paths, and the difference between them stabilizes.

Welfare. I measure the welfare effect of a policy on the household in country j by the equivalent variation in income: the proportional change in income along the laissez-faire path, in every period and at that path's prices, that gives the household the same lifetime utility as the policy path. Formally, let $C_{j,t}^{\mathrm{LF}}(E)$ denote the consumption aggregate that the household in country j chooses in period t of the laissez-faire path when its income is E, at that path's prices, and let $C_{j,t}$ denote its consumption aggregate on the policy path. The welfare gain EV_j solves

$$\sum_{t=t_0}^{\infty} \frac{1}{(1+\rho)^{\,t-t_0}}\, u\Big(C_{j,t}^{\mathrm{LF}}\big((1+\mathrm{EV}_j)\,E_{j,t}^{\mathrm{LF}}\big)\Big) = \sum_{t=t_0}^{\infty} \frac{1}{(1+\rho)^{\,t-t_0}}\, u\big(C_{j,t}\big), \tag{50}$$

where u is the period utility in (13), with log utility ($\sigma = 1$) as in the calibration, and $E_{j,t}^{\mathrm{LF}}$ is income on the laissez-faire path. I report EV_j in percent. Lifetime utility is evaluated over 400 periods and closed with a stationary continuation (Appendix D.1); the sign of every welfare number reported below is the same when the sums are truncated at 80 or 120 periods.

To see where the gain comes from, I decompose it by source. Under log utility, the period utility of a household is, up to a constant, the logarithm of its expenditure in excess of the cost of subsistence consumption, $E_{j,t} - P_{j,t}^{a}\,\bar{a}\,L_j$, less ϕ times the logarithm of the agricultural price index $P_{j,t}^{a}$. The difference in period utility between the policy path and the laissez-faire path therefore splits exactly into the change in the price index and the changes in the components of income in (40) and in the cost of subsistence consumption, each divided by the logarithmic mean of the two paths' excess expenditure (Appendix D.1). I report the contribution of each part in percentage points of EV_j, allocating the gain across the parts in proportion to their discounted sums. For Southern households, the parts are the food price index, the cost of subsistence

Table 3: The counterfactuals for Tanzania: the innovation intensity of type-S inputs and the subsidy rate in 2017 (the rate as a share of the mean R&D fixed cost of a product line of type-S inputs), the cost of the R&D subsidy in 2017 in percent of Southern GDP, household income before the tax, the rise in agricultural labor productivity in Tanzania relative to the US against the laissez-faire path, in percent, and the welfare gain of Southern households under the barrier interpretation of the wage gap and under the effective-labor interpretation. The subsidies are those of Sections 6.2 and 6.3, the stronger enforcement of intellectual property rights that of Section 6.4, and the other rows are those of Appendix D.4. The rises under the effective-labor interpretation differ from those shown by at most 1.8 percentage points.

			Cost	Rise in APD_t (%)		Welfare gain EV_S (%)	
Counterfactual	m_{S,t_0}	s_{t_0}/c_S	(% of GDP)	2037	2047	Barrier	Effective labor
Financed by Southern households							
Subsidy (ratio held)	0.28	0.36	0.33	9.6	11.6	2.40	0.10
Subsidy (all lines)	1.00	1.41	4.53	44.2	52.3	6.76	−3.23
Reforms without a subsidy							
Stronger IP enforcement	0.10	0	0	0.9	0.4	0.10	−0.06
Lower adaptation cost	0.11	0	0	1.2	0.6	0.21	0.03
Both reforms	0.20	0	0	3.9	2.5	0.53	0.01
Financed by Northern households							
Subsidy (ratio held)	0.28	0.36	0.33	9.6	11.6	2.67	0.59
Subsidy (all lines)	1.00	1.41	4.55	44.0	52.0	9.80	2.07

consumption, labor income, land income, the profits of imitators, and the tax that finances the subsidy. Labor income is the wages paid in each sector. By (41), workers in agriculture earn a fraction $1/(1+\tau_{w,S})$ of the non-agricultural wage, so a worker who leaves agriculture gains the whole measured wage gap. This is the *barrier interpretation* of the wage gap: the gap is a barrier to leaving agriculture, and labor income rises with every worker whom the policy moves out of agriculture. I return to the alternative interpretation below.

Table 3 summarizes the counterfactuals of this section, labeling the subsidy that holds the quality ratio "Subsidy (ratio held)" and the larger subsidy of Section 6.3, under which every product line of type-S inputs is improved, "Subsidy (all lines)", and Table 4 decomposes the welfare gains of Southern households. The subsidy raises the welfare of Southern households by 2.40 percent of their expenditure, even though they finance it. Of the 2.40 percentage points, labor income contributes 2.24. The mechanism is what Schultz (1953) characterized as the food problem: when agricultural productivity is low, a large share of the labor force must work in agriculture to meet the country's basic food needs (Restuccia et al., 2008). The subsidy raises Southern agricultural productivity by providing more appropriate inputs, so fewer workers are needed to meet subsistence consumption, and labor moves to non-agriculture: the farm employment share is 47.4 percent in 2037, against 50.1 percent under laissez-faire, so the sub-

Table 4: The welfare gain of Southern households and its sources, Tanzania: the gain EV_S in percent and the contributions, in percentage points, of the food price index, the cost of subsistence consumption, labor income, land income, the profits of imitators, and the tax that finances the subsidy. Southern households finance the subsidies. The contributions sum to the gain before rounding.

	EV_S	Food price index	Subsistence cost	Labor income	Land income	Imitators' profits	Tax
Subsidy (ratio held)							
Barrier interpretation	2.40	0.13	0.37	2.24	−0.05	0.01	−0.30
Effective-labor interpretation	0.10	0.13	0.53	0.00	−0.08	0.00	−0.49
Subsidy (all lines)							
Barrier interpretation	6.76	0.52	1.26	8.57	−0.18	0.03	−3.44
Effective-labor interpretation	−3.23	0.50	1.82	0.00	−0.29	0.02	−5.28

sidy moves 2.7 percent of the labor force out of agriculture, and each of these workers gains the wage gap.[31] The subsidy accelerates structural change rather than changing the allocation of labor in the long run: once subsistence consumption no longer matters, the farm employment share is nearly the same on both paths, whereas the rise in relative agricultural productivity is permanent. The cost of subsistence consumption contributes 0.37 points and the food price index 0.13 points. Both reflect the lower price of agricultural goods in the South: the first is the saving on the subsistence quantity $\bar{a}L_S$, and the second is the gain on expenditure in excess of the cost of subsistence consumption, of which agricultural goods account for the share ϕ. Land income contributes −0.05 points, because agricultural revenue falls, the profits of imitators 0.01 points, and the tax that finances the subsidy −0.30 points. The welfare gain thus comes from cheaper food and, above all, from the higher labor income of workers drawn out of agriculture by the subsidy. Northern households receive the subsidy as income of their R&D firm and pay the researchers of the additional product lines; their welfare rises by 0.0009 percent.

The effective-labor interpretation of the wage gap. The welfare gain depends on how the wage gap between agriculture and non-agriculture is interpreted. Under the barrier interpretation, the whole measured gap is a barrier to leaving agriculture. If instead the gap reflects differences in workers' effective labor, so that a worker in agriculture supplies fewer units of

[31] The relationship between low agricultural productivity and a large share of employment in agriculture need not hold if a country can import its food (Restuccia et al., 2008). Matsuyama (1992) shows that, with subsistence consumption of food, a rise in agricultural productivity releases labor from agriculture in a closed economy, whereas a more productive agricultural sector attracts labor in a small open economy. In the quantitative model, trade costs keep the Southern market for agricultural goods nearly closed: the calibration targets a share of 0.966 of Southern expenditure on agricultural goods allocated to the Southern variety (Section 5.1.1). Without trade costs, as in the baseline model, an improvement in type-S inputs instead raises Southern agricultural revenue and employment relative to the North's (Section 4.3.1).

effective labor than other workers and every unit earns the same wage in both sectors, then leaving agriculture yields no income gain. Under this *effective-labor interpretation*, I hold each household's labor income at its level in the laissez-faire equilibrium of 2017 along every path and recompute the equilibrium paths; the farm's demand for labor, prices, land income, the profits of imitators, and the tax adjust as before, and the labor part of the decomposition is zero by construction (Appendix D.1). The productivity results do not depend on the interpretation: the dynamic amplification factor on the recomputed paths differs from that of Section 6.1 by at most 1.21 percent, the subsidy raises Tanzania's relative agricultural productivity by 9.8 percent in 2037 and 11.9 percent in 2047, and it moves nearly the same workers out of agriculture, within 1.58 percentage points of the labor force. The welfare gain falls to 0.10 percent but remains positive (Table 4). The cost of subsistence consumption contributes 0.53 points and the food price index 0.13 points, land income -0.08 points, and the tax -0.49 points. The tax costs more than under the barrier interpretation because the same cost is set against a smaller income, and because nothing in Southern income rises with the policy to offset it. Gollin et al. (2014a) find that differences in hours worked and in human capital per worker account for roughly one third of the gap in value added per worker between the sectors on average, and one half in developing countries, and that a large gap remains. The model has neither hours nor human capital nor heterogeneous workers, so the barrier interpretation reads the whole measured gap as a barrier and the effective-labor interpretation reads none of it as one. The two interpretations bracket the gains of the workers who move, and a result that holds under both does not depend on the interpretation of labor reallocation.

The cost of the subsidy over time. Figure 4 shows the subsidy rate and the cost of the R&D subsidy in percent of Southern GDP from 2017 to 2082, for the subsidy that holds the quality ratio and for the larger subsidy of Section 6.3, under both interpretations of the wage gap. The rate rises under both programs: at a given quality ratio, the operating profit of a type-S product line falls as structural change contracts agriculture, so a larger subsidy is needed to induce the same innovation, and under the larger subsidy the margin rate also falls as the quality ratio rises and the scope for catch-up shrinks (Section 6.3). The cost as a share of GDP nevertheless falls under the barrier interpretation, from 0.33 percent in 2017 to 0.29 percent in 2082 under the subsidy that holds the quality ratio and from 4.53 to 3.00 percent under the larger subsidy, because Southern income grows as workers leave agriculture and earn the non-agricultural wage. Under the effective-labor interpretation, labor income is constant, and the share rises with the rate, to 0.50 and 5.39 percent. Over the whole horizon, the discounted cost of the subsidy that holds the quality ratio is 0.29 percent of discounted Southern GDP under the barrier interpretation and 0.49 percent under the effective-labor interpretation, against 3.06 and 5.30 percent for the larger subsidy; the tax rows of Table 4 value these costs in the welfare metric.

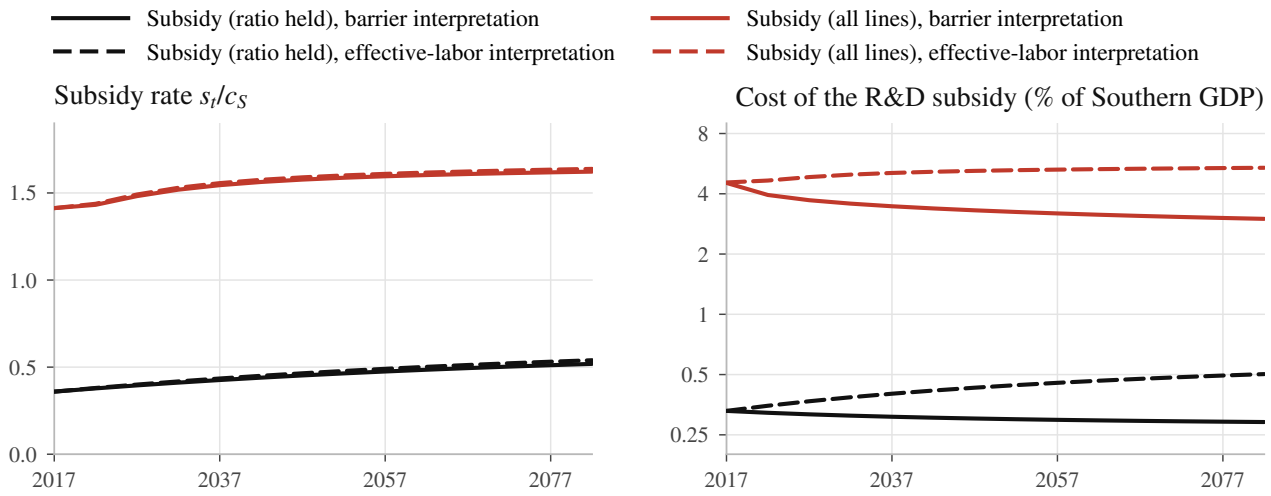


Figure 4: The subsidy rate and the cost of the R&D subsidy, Tanzania, 2017–2082: the subsidy that holds the quality ratio at its 2017 value and the subsidy under which every type-S line is improved, under the barrier interpretation of the wage gap (solid) and the effective-labor interpretation (dashed). The rate is the subsidy per improved product line as a share of the mean R&D fixed cost of a type-S line. The cost is the subsidy's total payment, $s_t m_{S,t}$, in percent of Southern GDP, household income before the tax, on the equilibrium path under the subsidy, which Southern households finance. Logarithmic axis in the right panel.

6.3 A Larger R&D Subsidy

The second program is a larger subsidy under which every product line of type-S inputs is improved from 2017 onward. In every period, the government sets the lowest subsidy rate at which $m_{S,t} = 1$, that is, $s_t = \bar{c}_S - \pi_{S,t}(\lambda_{t-1})$, so that the manager of the costliest product line is just willing to improve it. Figure 5 shows the paths. The subsidy raises the quality ratio toward the transferability ceiling instead of holding it (0.60 in 2047 and 0.65 in 2082), and the effective quality ratio from 0.31 to 0.62 and 0.66 (middle panel). The subsidy rate in 2017 equals 1.41 of the mean R&D fixed cost, about four times the rate of the headline program, and it rises to 1.62 by 2082: the rate must cover the costliest product line, the margin rate of type-S innovation falls as the quality ratio rises and the scope for catch-up shrinks (Section 2.3.2), and structural change contracts Southern agriculture, which the rising adoption of type-S inputs offsets only in part. The cost of the subsidy in 2017 is 4.53 percent of Southern GDP, or 30.8 percent of Southern agricultural revenue, 14 times the cost of the headline program.

Tanzania's relative agricultural productivity rises to 1/118 in 2037 and 1/112 in 2047, a rise of 44.2 and 52.3 percent, 4.6 times the rise under the headline program in 2037. By 2082, the path is within 0.1 percent of the undirected path of Section 6.1: a subsidy that improves every product line from 2017 onward removes the amplification over the following decades, but not the productivity lost in the meantime.

The welfare gain of Southern households is 6.76 percent under the barrier interpretation,

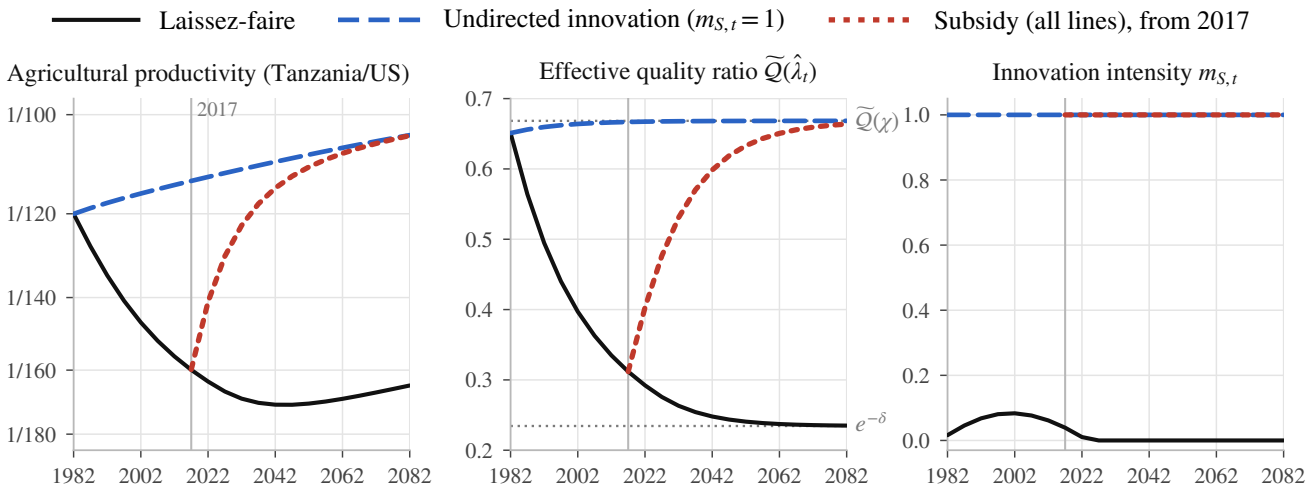


Figure 5: The larger R&D subsidy, under which every product line of type-S inputs is improved: the laissez-faire path and the undirected path from 1982, as in Figure 2, and the path under the subsidy from 2017, Tanzania and the US. The panels are those of Figure 2. Southern households finance the subsidy.

larger than under the headline program although the tax now contributes -3.44 points (Table 4). The subsidy moves 10.8 percent of the labor force out of agriculture by 2037, as the farm employment share falls to 39.3 percent against 50.1 percent under laissez-faire, and labor income contributes 8.57 points. Under the effective-labor interpretation, however, the program lowers Southern welfare by 3.23 percent. Without the gains of the workers who move, what the program delivers through the cost of subsistence consumption (1.82 points), the food price index (0.50), land income (-0.29), and the profits of imitators (0.02) does not cover the tax (-5.28 points). The sign of the larger program's welfare effect therefore depends on the interpretation of the wage gap, whereas the headline program raises Southern welfare under both. Both programs induce innovation that is privately unprofitable in every period, since no product line of type-S inputs is improved under laissez-faire from 2027 onward, and their benefits accrue in later periods through the quality state they leave behind: the quality ratio held at 0.24 under the headline program, and raised toward χ under the larger one.

6.4 Stronger Enforcement of Intellectual Property Rights and Other Counterfactuals

The last counterfactual is a reform rather than a subsidy: full enforcement of intellectual property rights in the South, $\zeta_S = 1$, from 2017 onward. The R&D firm then keeps the whole operating profit from sales to Southern farmers (Section 4.1.2), so the operating profit of a product line of type-S inputs rises. Figure 6 compares the laissez-faire path under the reform with the laissez-faire path without it. The innovation intensity in 2017 rises from 0.04 to 0.10, but it stays

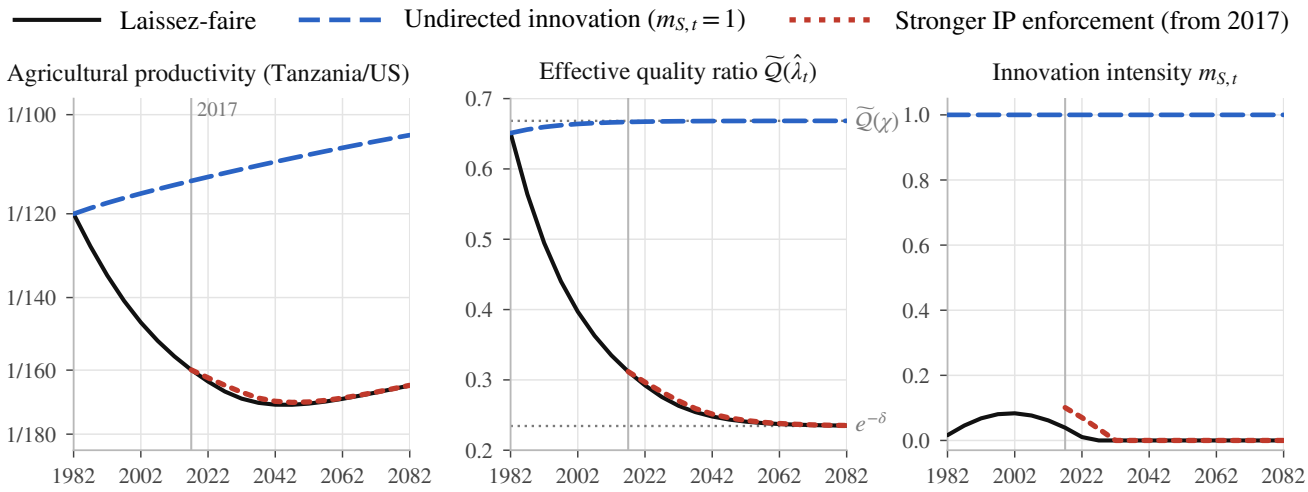


Figure 6: Stronger enforcement of intellectual property rights in the South ($\zeta_S = 1$ from 2017): the laissez-faire path without the reform and the undirected path from 1982, as in Figure 2, and the laissez-faire path under the reform from 2017, Tanzania and the US. The panels are those of Figure 2.

below $\overline{m}(\lambda_{t-1})$, so the quality ratio keeps falling (0.10 in 2047 against 0.09 without the reform), and the intensity reaches zero in 2032, one period later than without the reform. The reform does not alter the long-run trajectory of the economy. The reform raises Tanzania's relative agricultural productivity by only 0.9 percent in 2037 and 0.4 percent in 2047, and the two paths converge thereafter.

Southern households gain 0.10 percent under the barrier interpretation (Table 3). The reform costs nothing, so the number is the same whichever households would have financed a subsidy. Its effect on Southern welfare is subtle. Under weak enforcement, imitators in the South appropriate part of the operating profit from sales to Southern farmers, and their profits, $\text{Profit}^{\chi}_{S,t}$, are Southern income. Full enforcement transfers these profits to the Northern patent holders, so Southern income falls on this account. What works in the other direction is the higher innovation intensity: slightly cheaper agricultural goods, slightly lower agricultural revenue, and a few workers moving out of agriculture. Under the barrier interpretation, the gains of these workers outweigh the lost profits of imitators. Under the effective-labor interpretation, they are zero, and the reform lowers Southern welfare by 0.06 percent.

Appendix D.4 presents further counterfactuals, with their main results in Table 3. In a scenario in which the cost of adapting inputs to Southern conditions equals the cost of improving type-N inputs ($f = 0$), the innovation intensity in 2017 rises to 0.11, Tanzania's relative agricultural productivity rises by 1.2 and 0.6 percent in 2037 and 2047, and Southern welfare rises by 0.21 percent. With both reforms together, the intensity rises to 0.20, relative agricultural productivity rises by 3.9 and 2.5 percent, and Southern welfare rises by 0.53 percent. Neither scenario changes the long-run trajectory of the economy. Combining the headline subsidy with a reform

changes its welfare gain little, to 2.37 percent with full enforcement and 2.50 percent with the lower adaptation cost, against 2.40 percent, since the subsidy holds the innovation intensity at the same target with or without the reform. Finally, when Northern households finance the subsidies, Southern households gain 2.67 percent under the headline program and 9.80 percent under the larger one, and the welfare of Northern households falls by less than 0.03 percent, the cost of the additional research.

7 Conclusion

Agricultural productivity differences across countries are far larger than non-agricultural ones, and the agricultural technology available to poor countries is often inappropriate for their environments. This paper explains how this inappropriateness emerges and persists in a two-country dynamic general equilibrium model in which agricultural technology is embodied in intermediate inputs adapted to Northern or Southern conditions and R&D firms direct innovation toward the input type that is profitable to improve. The key mechanism is dynamic lock-in through directed innovation. When the market for the inputs adapted to the South is too small for operating profits to cover the fixed cost of R&D, these inputs are not improved and fall further behind as the frontier improves, which shrinks their market further, until Southern farmers use the frontier inputs despite their environmental mismatch. Directed innovation thus locks poor countries into inappropriate technology and amplifies agricultural productivity differences over time. This lock-in is constrained inefficient when the future gains from improving the Southern inputs cover the cost of R&D: a social planner who takes the within-period equilibrium structure as given then improves them where private firms do not, because these gains accrue to households and future innovators, and the innovator has no claim on them that it can sell.

Quantitatively, calibrating the model to Tanzania and the US, with the key parameters disciplined by the micro estimates of Moscona and Sastry (2025) and by the adoption of modern varieties during the Green Revolution, I find that directed innovation since 1982 amplifies agricultural productivity differences between Tanzania and the US by 42 percent in 2017 and 51 percent in the long run, relative to a path on which both input types improve in every period. An R&D subsidy introduced in 2017 that keeps the inputs adapted to Tanzania from falling further behind the frontier raises Tanzania's agricultural labor productivity relative to the US by 12 percent by 2047 and raises the welfare of its households by 2.4 percent, even though they finance it. The gain comes largely from the higher labor income of the workers whom the subsidy moves out of agriculture, and it remains positive when the sectoral wage gap is read as differences in workers' effective labor rather than as a barrier to leaving agriculture. A larger subsidy narrows the difference further, but whether it raises welfare depends on this reading, and stronger enforcement of intellectual property rights alone does not alter the economy's long-run trajectory. Both subsidies pay for innovation that is privately unprofitable but leaves

behind an improved quality state whose benefits accrue in later periods, the gains that private firms do not internalize. The inappropriateness of the technology available to poor countries is thus not a fixed feature of their environments but an outcome of the direction of innovation, and R&D policies that redirect innovation toward technologies adapted to their conditions can narrow agricultural productivity differences and raise welfare in poor countries, even when they finance these policies themselves.

References

Acemoglu, D., Aghion, P., Barrage, L., & Hémous, D. (2024). Climate Change, Directed Innovation, and Energy Transition: The Long-Run Consequences of the Shale Gas Revolution. *Mimeo*.

Acemoglu, D., & Zilibotti, F. (2001). Productivity Differences. *The Quarterly Journal of Economics*, *116*(2), 563–606.

Adamopoulos, T. (2011). Transportation Costs, Agricultural Productivity, and Cross-Country Income Differences. *International Economic Review*, *52*(2), 489–521.

Adamopoulos, T., Brandt, L., Leight, J., & Restuccia, D. (2022). Misallocation, Selection, and Productivity: A Quantitative Analysis with Panel Data from China. *Econometrica*, *90*(3), 1261–1282.

Adamopoulos, T., & Restuccia, D. (2014). The Size Distribution of Farms and International Productivity Differences. *American Economic Review*, *104*(6), 1667–1697.

Adamopoulos, T., & Restuccia, D. (2022). Geography and Agricultural Productivity: Cross-Country Evidence from Micro Plot-Level Data. *Review of Economic Studies*, *89*(4), 1629–1653.

AgbioInvestor. (2024). *Time and Cost of New Agrochemical Product Discovery, Development and Registration.* Study results, a study on behalf of CropLife International.

AgbioInvestor. (2026). *Time and Cost of New Agrochemical Product Discovery, Development and Registration.* Study results, version 2, a study on behalf of CropLife International.

Aghion, P., & Howitt, P. (1992). A Model of Growth through Creative Destruction. *Econometrica*, *60*(2), 323–351.

Akerman, A., Moscona, J., Pellegrina, H. S., & Sastry, K. A. (2026). Public R&D Meets Economic Development: Embrapa and Brazil's Agricultural Revolution. *Mimeo*.

Alvarez-Cuadrado, F., & Poschke, M. (2011). Structural Change Out of Agriculture: Labor Push versus Labor Pull. *American Economic Journal: Macroeconomics*, *3*(3), 127–158.

Arrow, K. J. (1962). Economic Welfare and the Allocation of Resources for Invention. In R. R. Nelson (Ed.), *The rate and direction of inventive activity: Economic and social factors* (pp. 609–626). Princeton, NJ: Princeton University Press.

Atkinson, A. B., & Stiglitz, J. E. (1969). A New View of Technological Change. *The Economic Journal*, *79*(315), 573–578.

Bandiera, O., & Rasul, I. (2006). Social Networks and Technology Adoption in Northern Mozambique. *The Economic Journal*, *116*(514), 869–902.

Bartelme, D., Costinot, A., Donaldson, D., & Rodríguez-Clare, A. (2025). The Textbook Case for Industrial Policy: Theory Meets Data. *Journal of Political Economy*, *133*(5), 1527–1573.

Basu, S., & Weil, D. N. (1998). Appropriate Technology and Growth. *The Quarterly Journal of Economics*, *113*(4), 1025–1054.

Beaman, L., BenYishay, A., Magruder, J., & Mobarak, A. M. (2021). Can Network Theory-Based

Targeting Increase Technology Adoption? *American Economic Review*, *111*(6), 1918–1943.

BenYishay, A., & Mobarak, A. M. (2019). Social Learning and Incentives for Experimentation and Communication. *Review of Economic Studies*, *86*(3), 976–1009.

Bianchi, J. (2011). Overborrowing and Systemic Externalities in the Business Cycle. *American Economic Review*, *101*(7), 3400–3426.

Bold, T., Kaizzi, K. C., Svensson, J., & Yanagizawa-Drott, D. (2017). Lemon Technologies and Adoption: Measurement, Theory and Evidence from Agricultural Markets in Uganda. *The Quarterly Journal of Economics*, *132*(3), 1055–1100.

Boppart, T. (2014). Structural Change and the Kaldor Facts in a Growth Model With Relative Price Effects and Non-Gorman Preferences. *Econometrica*, *82*(6), 2167–2196.

Boppart, T., Kiernan, P., Krusell, P., & Malmberg, H. (2026). The Macroeconomics of Intensive Agriculture. *Mimeo*.

Brooks, W., & Donovan, K. (2026). *Industrial Policy with Development Characteristics: Fertilizer Policy in Times of Crisis* (NBER Working Paper No. 35126). National Bureau of Economic Research.

Bustos, P., Caprettini, B., & Ponticelli, J. (2016). Agricultural Productivity and Structural Transformation: Evidence from Brazil. *American Economic Review*, *106*(6), 1320–1365.

Campi, M., & Nuvolari, A. (2021). Intellectual Property Rights and Agricultural Development: Evidence from a Worldwide Index of IPRs in Agriculture (1961–2018). *The Journal of Development Studies*, *57*(4), 650–668.

Caselli, F. (2005). Accounting for Cross-Country Income Differences. In P. Aghion & S. N. Durlauf (Eds.), *Handbook of economic growth* (Vol. 1A, pp. 679–741). Amsterdam: Elsevier.

Caselli, F., & Coleman, W. J. (2006). The World Technology Frontier. *American Economic Review*, *96*(3), 499–522.

Caunedo, J., Felix, M., & Manysheva, K. (2026). Constructing Structural Change. *Mimeo*.

Caunedo, J., & Keller, E. (2021). Capital Obsolescence and Agricultural Productivity. *The Quarterly Journal of Economics*, *136*(1), 505–561.

Chen, C. (2017). Untitled Land, Occupational Choice, and Agricultural Productivity. *American Economic Journal: Macroeconomics*, *9*(4), 91–121.

Chen, C. (2020). Technology Adoption, Capital Deepening, and International Productivity Differences. *Journal of Development Economics*, *143*, 102388.

Choi, J., & Levchenko, A. A. (2025). The Long-Term Effects of Industrial Policy. *Journal of Monetary Economics*, *152*, 103779.

Comin, D., Lashkari, D., & Mestieri, M. (2021). Structural Change with Long-Run Income and Price Effects. *Econometrica*, *89*(1), 311–374.

Conley, T. G., & Udry, C. R. (2010). Learning about a New Technology: Pineapple in Ghana. *American Economic Review*, *100*(1), 35–69.

Costinot, A., Donaldson, D., & Smith, C. (2016). Evolving Comparative Advantage and the Im-

pact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World. *Journal of Political Economy*, *124*(1), 205–248.

Dávila, E., & Korinek, A. (2018). Pecuniary Externalities in Economies with Financial Frictions. *Review of Economic Studies*, *85*(1), 352–395.

Dercon, S., & Christiaensen, L. (2011). Consumption Risk, Technology Adoption and Poverty Traps: Evidence from Ethiopia. *Journal of Development Economics*, *96*(2), 159–173.

Donovan, K. (2021). The Equilibrium Impact of Agricultural Risk on Intermediate Inputs and Aggregate Productivity. *Review of Economic Studies*, *88*(5), 2275–2307.

Duflo, E., Kremer, M., & Robinson, J. (2011). Nudging Farmers to Use Fertilizer: Theory and Experimental Evidence from Kenya. *American Economic Review*, *101*(6), 2350–2390.

Emerick, K., de Janvry, A., Sadoulet, E., & Dar, M. H. (2016). Technological Innovations, Downside Risk, and the Modernization of Agriculture. *American Economic Review*, *106*(6), 1537–1561.

Evenson, R. E., & Gollin, D. (Eds.). (2003). *Crop Variety Improvement and Its Effect on Productivity: The Impact of International Agricultural Research*. Wallingford, UK: CABI Publishing.

Foster, A. D., & Rosenzweig, M. R. (1995). Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *Journal of Political Economy*, *103*(6), 1176–1209.

Foster, A. D., & Rosenzweig, M. R. (2010). Microeconomics of Technology Adoption. *Annual Review of Economics*, *2*, 395–424.

Fujiwara, I., & Matsuyama, K. (2024). A Technology-Gap Model of 'Premature' Deindustrialization. *American Economic Review*, *114*(11), 3714–3745.

Geanakoplos, J. D., & Polemarchakis, H. M. (1986). Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market Is Incomplete. In W. P. Heller, R. M. Starr, & D. A. Starrett (Eds.), *Uncertainty, information, and communication: Essays in honor of Kenneth J. Arrow* (Vol. 3, pp. 65–95). Cambridge: Cambridge University Press.

Gerschenkron, A. (1962). *Economic Backwardness in Historical Perspective: A Book of Essays*. Cambridge, MA: Belknap Press of Harvard University Press.

Gollin, D., Lagakos, D., & Waugh, M. E. (2014a). The Agricultural Productivity Gap. *The Quarterly Journal of Economics*, *129*(2), 939–993.

Gollin, D., Lagakos, D., & Waugh, M. E. (2014b). Agricultural Productivity Differences across Countries. *American Economic Review: Papers and Proceedings*, *104*(5), 165–170.

Gollin, D., Parente, S. L., & Rogerson, R. (2002). The Role of Agriculture in Development. *American Economic Review*, *92*(2), 160–164.

Gollin, D., Parente, S. L., & Rogerson, R. (2007). The Food Problem and the Evolution of International Income Levels. *Journal of Monetary Economics*, *54*(4), 1230–1255.

Gottlieb, C., & Grobovšek, J. (2019). Communal Land and Agricultural Productivity. *Journal of Development Economics*, *138*, 135–152.

Griliches, Z. (1957). Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica*, *25*(4), 501–522.

Grossman, G. M., & Helpman, E. (1991). Quality Ladders in the Theory of Growth. *Review of Economic Studies*, *58*(1), 43–61.

Harrison, A., & Rodríguez-Clare, A. (2010). Trade, Foreign Investment, and Industrial Policy for Developing Countries. In D. Rodrik & M. Rosenzweig (Eds.), *Handbook of development economics* (Vol. 5, pp. 4039–4214). Amsterdam: Elsevier.

Hausmann, R., & Rodrik, D. (2003). Economic Development as Self-Discovery. *Journal of Development Economics*, *72*(2), 603–633.

Herrendorf, B., Rogerson, R., & Valentinyi, Á. (2013). Two Perspectives on Preferences and Structural Transformation. *American Economic Review*, *103*(7), 2752–2789.

Herrendorf, B., & Schoellman, T. (2018). Wages, Human Capital, and Barriers to Structural Transformation. *American Economic Journal: Macroeconomics*, *10*(2), 1–23.

Inklaar, R., Marapin, R., & Gräler, K. (2023). *Tradability and Sectoral Productivity Differences across Countries* (GGDC Research Memorandum No. 195). Groningen Growth and Development Centre, University of Groningen.

Itskhoki, O., & Moll, B. (2019). Optimal Development Policies With Financial Frictions. *Econometrica*, *87*(1), 139–173.

Jeanne, O., & Korinek, A. (2019). Managing Credit Booms and Busts: A Pigouvian Taxation Approach. *Journal of Monetary Economics*, *107*, 2–17.

Juhász, R. (2018). Temporary Protection and Technology Adoption: Evidence from the Napoleonic Blockade. *American Economic Review*, *108*(11), 3339–3376.

Juhász, R., Lane, N., & Rodrik, D. (2024). The New Economics of Industrial Policy. *Annual Review of Economics*, *16*, 213–242.

Karlan, D., Osei, R., Osei-Akoto, I., & Udry, C. (2014). Agricultural Decisions after Relaxing Credit and Risk Constraints. *The Quarterly Journal of Economics*, *129*(2), 597–652.

Kim, M., Lee, M., & Shin, Y. (2026). The Plant-Level View of an Industrial Policy: The Korean Heavy Industry Drive of 1973. *Mimeo*.

Kremer, M., & Glennerster, R. (2004). *Strong Medicine: Creating Incentives for Pharmaceutical Research on Neglected Diseases*. Princeton, NJ: Princeton University Press.

Kremer, M., Levin, J., & Snyder, C. M. (2020). Advance Market Commitments: Insights from Theory and Experience. *AEA Papers and Proceedings*, *110*, 269–273.

Lagakos, D., & Schoellman, T. (2026). *Accounting for Cross-Country Income Differences Revisited* (NBER Working Paper No. 35826). National Bureau of Economic Research.

Lagakos, D., & Waugh, M. E. (2013). Selection, Agriculture, and Cross-Country Productivity Differences. *American Economic Review*, *103*(2), 948–980.

Lane, N. (2025). Manufacturing Revolutions: Industrial Policy and Industrialization in South Korea. *The Quarterly Journal of Economics*, *140*(3), 1683–1741.

Lashkaripour, A., & Lugovskyy, V. (2023). Profits, Scale Economies, and the Gains from Trade

and Industrial Policy. *American Economic Review*, *113*(10), 2759–2808.

Liu, E. (2019). Industrial Policies in Production Networks. *The Quarterly Journal of Economics*, *134*(4), 1883–1948.

Mabaya, E., Hassena, M., Waithaka, M., Mugoya, M., Camara, M., Tihanyi, K., … Damba, B. (2021). *Ethiopia Country Report 2021.* Country report, The African Seed Access Index (TASAI).

Mabaya, E., Kachule, R., Waithaka, M., Mugoya, M., Damba, B., Camara, M., & Tihanyi, K. (2022). *Malawi Country Report 2022.* Country report, The African Seed Access Index (TASAI).

Mabaya, E., Kyotalimye Wandukwa, M., Ssebuliba, R., Waithaka, M., Mugoya, M., Anyoso, W., … Tihanyi, K. (2025). *Uganda Country Report 2025.* Country report, The African Seed Access Index (TASAI).

Mabaya, E., Mizambwa, F., Komba, C., Waithaka, M., Mugoya, M., Tihanyi, K., & Kanyenji, G. (2021). *Tanzania Country Report 2021.* Country report, The African Seed Access Index (TASAI).

Mabaya, E., Urinzwenimana, C., Ngerero, N. G., Waithaka, M., Mugoya, M., Tihanyi, K., & Kanyenji, G. (2021). *Rwanda Country Report 2021.* Country report, The African Seed Access Index (TASAI).

Mabaya, E., Waithaka, M., Mugoya, M., Kanyenji, G., Tihanyi, K., Ssebuliba, R., & Kyotalimye, M. (2020). *Uganda Country Report 2020.* Country report, The African Seed Access Index (TASAI).

Manysheva, K. (2024). Land Property Rights, Financial Frictions, and Resource Allocation in Developing Countries. *Mimeo*.

Matsuyama, K. (1992). Agricultural Productivity, Comparative Advantage, and Economic Growth. *Journal of Economic Theory*, *58*(2), 317–334.

Mburu, J., Mabaya, E., Waithaka, M., Mugoya, M., Damba, B., Camara, M., & Tihanyi, K. (2022). *Kenya Country Report 2022.* Country report, The African Seed Access Index (TASAI).

Moscona, J., & Sastry, K. A. (2025). Inappropriate Technology: Evidence from Global Agriculture. *Mimeo*.

Phillips McDougall. (2011). *The Cost and Time Involved in the Discovery, Development and Authorisation of a New Plant Biotechnology Derived Trait* (R&D Study). Pathhead, Midlothian, United Kingdom: A consultancy study for CropLife International; Phillips McDougall.

Phillips McDougall. (2016). *The Cost of New Agrochemical Product Discovery, Development and Registration in 1995, 2000, 2005–8 and 2010–2014: R&D Expenditure in 2014 and Expectations for 2019* (R&D Study). Pathhead, Midlothian, United Kingdom: A consultancy study for CropLife International, CropLife America and the European Crop Protection Association; Phillips McDougall.

Porzio, T. (2017). Cross-Country Differences in the Optimal Allocation of Talent and Technology. *Mimeo*.

Restuccia, D., Yang, D. T., & Zhu, X. (2008). Agriculture and Aggregate Productivity: A Quantitative Cross-Country Analysis. *Journal of Monetary Economics*, *55*(2), 234–250.

Rossi, F. (2022). The Relative Efficiency of Skilled Labor across Countries: Measurement and Interpretation. *American Economic Review*, *112*(1), 235–266.

Schultz, T. W. (1953). *The Economic Organization of Agriculture*. New York: McGraw-Hill.

Stewart, F. (1978). *Technology and Underdevelopment*. London: Macmillan.

Stiglitz, J. E. (1982). The Inefficiency of the Stock Market Equilibrium. *Review of Economic Studies*, *49*(2), 241–261.

Suri, T. (2011). Selection and Comparative Advantage in Technology Adoption. *Econometrica*, *79*(1), 159–209.

Suri, T., & Udry, C. (2022). Agricultural Technology in Africa. *Journal of Economic Perspectives*, *36*(1), 33–56.

Tombe, T. (2015). The Missing Food Problem: Trade, Agriculture, and International Productivity Differences. *American Economic Journal: Macroeconomics*, *7*(3), 226–258.

USDA Economic Research Service. (2026). *International Agricultural Productivity.* Data product, U.S. Department of Agriculture, Economic Research Service. Data and methods as of January 2026.

A Proofs and Derivations for the Baseline Model

This appendix contains the derivations and proofs for the baseline model. Section A.1 states the assumptions of the baseline model formally. Section A.2 derives farmers' input demand and the prices of intermediate inputs, and proves Lemma 1. Section A.3 characterizes the equilibrium of a given period in full. Section A.4 shows that this equilibrium is unique. Section A.5 proves Proposition 1 and derives the corresponding expression for agricultural employment shares. Section A.6 allows R&D to take place in both countries and shows that it occurs only in the North. Section A.7 characterizes the path of the quality ratio under laissez-faire, which Section 2.3.2 describes in words. Sections A.8–A.10 derive the constrained planner's problem of Section 3, prove its results, and introduce the R&D subsidy that implements the planner's decisions.

Throughout, $\psi = \alpha^2$, $\gamma := 1 - \alpha - \beta$, $\Gamma := 1 + \gamma(\varepsilon - 1)$, and $\mu = 1/2$. The innovation decision $m_{k,t} \in \{0,1\}$ equals one if type-k inputs are improved in period t and zero otherwise; Section 3 uses $m_{S,t}$ as the planner's decision variable, and Section 4 extends $m_{k,t}$ to the innovation intensity. To shorten the derivations, I use the following symbols, which do not appear in the main text. Let $R_{j,t} := p^a_{j,t} Y^a_{j,t}$ denote agricultural revenue in country j, $R^W_t := R_{N,t} + R_{S,t}$ world agricultural revenue, and $s_{S,t} := R_{S,t}/R^W_t$ the South's share of world agricultural revenue. Finally, let $h^*_{j,t} := h_{j,\tau^*_{j,t},t}$ denote the price-adjusted effective quality of the input type used in country j.

A.1 Assumptions

Section 2 introduces the assumptions of the baseline model where they are first used. This section collects them and states them formally. It also specifies the order of decisions within Stage 2, on which the pricing results rely.

Assumption A.1. *If $h_{j,N,t} = h_{j,S,t}$, the farmer in country j uses the input type adapted to the local environment, so that $\tau^*_{j,t} = j$.*

Assumption A.2.

(a) Disclosure at patent expiry. *When the patent on a vintage expires, its blueprint becomes publicly available. From the following period onward, the competitive fringe can produce this vintage at marginal cost ψ. Among the vintages of an input type whose patents have expired, only the most recent one is traded. If type-k innovation does not occur in period t, type-k inputs are sold at price ψ.*

(b) Order of decisions within Stage 2. *Stage 2 consists of three steps. In Step 1, farmers choose an input type and hire labor, anticipating the prices of intermediate inputs. In Step 2, the patent holder of each improved input type posts a price, and the competitive fringe offers the previous vintage at marginal cost ψ. In Step 3, farmers choose between the two vintages of the input type chosen in Step 1 and decide how much to purchase. A farmer who is indifferent between the two*

vintages buys from the patent holder. Production, trade, and consumption then take place. Farmers have rational expectations, so the anticipated and the posted prices coincide in equilibrium.

(c) Perceived demand. *When posting its price in Step 2, the patent holder takes as given agricultural output prices, household incomes, and the choices that farmers made in Step 1.*

Part (a) is the standard convention of quality-ladder models (Grossman & Helpman, 1991; Aghion & Howitt, 1992): once a patent expires, the vintage can be produced by anyone, and it becomes the alternative that limits the price of the next improvement. Part (b) reflects the different time scales of the two decisions. Which type of seed or input to use is decided before the season begins, and it depends on seed systems, extension services, and learning (Suri, 2011; Foster & Rosenzweig, 2010; Bold et al., 2017), whereas input purchases are made within the season. As a result, a patent holder cannot attract farmers who chose the other input type by lowering its price within the period. Part (c) is the usual assumption that a firm takes aggregate variables as given.

Assumption A.3.

(a) Non-drastic frontier innovation. $1+\xi<1/\alpha$.

(b) Frontier innovation in every period. $\Pi_{N,t}\geq 0$ *in every period t, whether or not type-S innovation occurs in that period.*

(c) Location of R&D. *All R&D firms are located in the North.*

(d) One-period R&D firms. *An R&D firm that innovates in period t earns the operating profit of its one-period patent and exits at the end of the period. It holds no claim on research opportunities in later periods and cannot contract with the R&D firms that enter later. Northern households own the R&D firms.*

Part (a) is the hypothesis of Lemma 1. Part (b) is the assumption introduced in Section 2.2 to focus on type-S innovation. Section A.6 derives part (c) from a comparison of R&D fixed costs across countries. Part (d) implies that the entry condition (12) compares the operating profit of a single period with the fixed cost of R&D: a firm that innovates in period t earns nothing from the higher quality on which later innovations build.

Assumption A.4.

(a) Parameters and initial condition. $\alpha,\beta\in(0,1)$ *with* $\alpha+\beta<1$; $\phi\in(0,1)$; $\varepsilon>1$; $\xi>0$; $\delta>0$; $a\in(0,1]$; $\chi\in(0,1]$; $\eta>0$; $\rho>0$; $\sigma>0$; A_j^a, A_j^n, *and* L_j *are strictly positive for* $j\in\{N,S\}$; *and the initial base qualities are strictly positive with* $\lambda_{-1}\in(0,\chi]$.

(b) Positive world income net of R&D fixed costs. $A_N^n L_N + A_S^n L_S > 2c$.

(c) Positive non-agricultural employment. *In every period,* $L_{j,t}^n>0$ *for* $j\in\{N,S\}$.

(d) Finite lifetime utility. $(1+\xi)^{\alpha\phi(1-\sigma)} < 1+\rho$.

Under part (a) and Assumption A.3(b), the quality ratio remains in $(0,\chi]$ in every period: by the law of motion (17), the map T sends $(0,\chi]$ into itself, and the quality ratio declines in periods without type-S innovation. Hence $q_{S,t-1} \leq \chi q_{N,t-1}$ in every period, which is the condition under which the law of motion for type-S inputs takes the form (10). Part (b) ensures that world agricultural revenue is strictly positive, whether or not type-S innovation occurs. Part (c) implies that wages are equalized across sectors. Part (d) ensures that lifetime utility (13) is finite. By (A.1) and Proposition A.1(iii), agricultural output prices in period t are proportional to $(q_{N,t-1})^{-\alpha}$. Consumption $C_{j,t}$ therefore equals $(q_{N,t-1})^{\alpha\phi}$ times a factor that depends only on λ_{t-1} and on whether type-S innovation occurs in period t, and this factor is bounded above and bounded away from zero. Because $q_{N,t-1}$ grows by the factor $1+\xi$ per period, the term $(q_{N,t-1})^{\alpha\phi}$ grows by the factor $(1+\xi)^{\alpha\phi}$ per period. Part (d) holds for every $\rho > 0$ if $\sigma \geq 1$.

Assumption A.5.

(a) $\lambda^D \leq e^{-\delta}$, *with strict inequality if* $a = 1$.

(b) $\chi > e^{-\delta}$.

Part (a) ensures that type-S innovation, whether drastic or non-drastic, does not change the input type that Southern farmers use in the period in which it occurs: Southern farmers use type-S inputs in period t if and only if $\lambda_{t-1} \geq e^{-\delta}$ (Proposition A.1(ii)). Part (b) ensures that there are quality ratios at which Southern farmers use type-S inputs. Neither Lemma 1 nor Proposition 1 requires Assumption A.5. In Sections A.1–A.6, it is used only where it is stated explicitly.

A.2 Input Demand and Pricing

Input type choice and input demand. Consider the problem of the representative farmer in country j in Section 2.1.1. The two input types enter the production function (1) only through the total number of effective units, $B_{j,N} q_{N,t} x_{j,N,t} + B_{j,S} q_{S,t} x_{j,S,t}$. One effective unit obtained from type-k inputs costs $p^x_{j,k,t}/(B_{j,k} q_{k,t}) = 1/h_{j,k,t}$. The farmer therefore uses only the input type with the highest price-adjusted effective quality $h_{j,k,t}$, which gives (4). If $h_{j,N,t} = h_{j,S,t}$, the farmer uses the input type adapted to the local environment (Assumption A.1).

Given the input type $\tau^*_{j,t}$, the first-order conditions for intermediate inputs and labor are

$$p^x_{j,\tau^*_{j,t},t}\, x_{j,\tau^*_{j,t},t} = \alpha R_{j,t} \qquad \text{and} \qquad w^a_{j,t} L^a_{j,t} = \beta R_{j,t},$$

which are the demands for intermediate inputs and labor in Section 2.1.1. The farmer spends the share α of agricultural revenue on intermediate inputs and the share β on labor, regardless of the input price. The remaining share is the farmer's profit, $\text{Profit}^a_{j,t} = \gamma R_{j,t}$, which is the income of the fixed factor.

The farmer uses $B_{j,\tau^*_{j,t}} q_{\tau^*_{j,t},t}\, x_{j,\tau^*_{j,t},t} = \alpha\, h^*_{j,t}\, R_{j,t}$ effective units of intermediate inputs. Substituting this expression and the labor demand into the production function (1) and multiplying by $p^a_{j,t}$ gives $R_{j,t} = p^a_{j,t} A^a_j \left(\beta R_{j,t}/w^a_{j,t}\right)^\beta \left(\alpha\, h^*_{j,t}\, R_{j,t}\right)^\alpha$, and hence

$$R^\gamma_{j,t} = \alpha^\alpha \beta^\beta A^a_j \left(w^a_{j,t}\right)^{-\beta} p^a_{j,t} \left(h^*_{j,t}\right)^\alpha. \tag{A.1}$$

Equation (A.1) relates agricultural revenue in country j to the agricultural output price, the wage, and the price-adjusted effective quality of the inputs in use.

Wages and agricultural labor productivity. The optimality condition of the non-agricultural firm is $w^n_{j,t} = A^n_j$. Because labor is mobile across sectors and non-agricultural employment is strictly positive (Assumption A.4(c)), wages are equalized across sectors: $w^a_{j,t} = w^n_{j,t} = A^n_j$ for $j \in \{N,S\}$ and $w^r_{N,t} = A^n_N$. The R&D fixed cost (11) is therefore $c = A^n_N/\eta$. The labor demand of the farmer then gives

$$L^a_{j,t} = \frac{\beta R_{j,t}}{A^n_j} \qquad \text{and} \qquad \frac{Y^a_{j,t}}{L^a_{j,t}} = \frac{A^n_j}{\beta\, p^a_{j,t}}, \tag{A.2}$$

where the second equality uses $Y^a_{j,t} = R_{j,t}/p^a_{j,t}$. Agricultural labor productivity in country j equals the wage divided by $\beta p^a_{j,t}$. Agricultural technology therefore affects agricultural labor productivity only through the agricultural output price.

The choice between vintages. Suppose that type-k innovation occurs in period t. In Step 3 of Stage 2, a farmer in country j who chose type-k inputs and labor L in Step 1 can buy the improved vintage, of base quality $q_{k,t}$, from the patent holder at price p, or the previous vintage, of base quality $q_{k,t-1}$, from the competitive fringe at price ψ. If the farmer buys x units of a vintage of base quality q at price p', the payoff is $p^a_{j,t} A^a_j L^\beta (B_{j,k}\, q\, x)^\alpha - p' x$. Maximizing over x gives the payoff

$$(1-\alpha)\,\alpha^{\frac{\alpha}{1-\alpha}} \left(p^a_{j,t} A^a_j L^\beta\right)^{\frac{1}{1-\alpha}} \left(\frac{B_{j,k}\, q}{p'}\right)^{\frac{\alpha}{1-\alpha}},$$

which is strictly increasing in q/p'. The farmer therefore buys the improved vintage if and only if

$$\frac{q_{k,t}}{p} \geq \frac{q_{k,t-1}}{\psi} \qquad \Longleftrightarrow \qquad p \leq G_{k,t}\, \psi, \tag{A.3}$$

where indifference is resolved in favor of the patent holder (Assumption A.2(b)). Condition (A.3) is the same for farmers in both countries, because the appropriateness $B_{j,k}$ is common to the two vintages of an input type.

Lemma A.1. *Suppose that Assumption A.2 holds. If type-k innovation occurs in period t, the patent holder posts the price*

$$p^x_{j,k,t} = \min\left\{\alpha,\ G_{k,t}\, \psi\right\}$$

in both countries $j \in \{N,S\}$, which is (7). If type-k innovation does not occur in period t, type-k inputs are sold at price ψ. The price of type-k inputs depends neither on the size of their market nor on which country's farmers use them.

Proof. Consider the patent holder of type-k inputs in Step 2 of Stage 2. By the first-order condition for x in Step 3, a farmer who chose type-k inputs in Step 1 buys a quantity proportional to $p^{-1/(1-\alpha)}$ from the patent holder if $p \leq G_{k,t}\,\psi$, and nothing otherwise, by (A.3). Under Assumption A.2(c), the patent holder therefore faces the demand $\Theta_{k,t}\,p^{-1/(1-\alpha)}$ for $p \leq G_{k,t}\,\psi$, where $\Theta_{k,t} \geq 0$ collects the terms that the patent holder takes as given, and zero demand for $p > G_{k,t}\,\psi$. For $p \leq G_{k,t}\,\psi$, its operating profit is $(p-\psi)\,\Theta_{k,t}\,p^{-1/(1-\alpha)}$, and

$$\frac{d}{dp}\left[(p-\psi)\,p^{-\frac{1}{1-\alpha}}\right] = \frac{\psi-\alpha p}{1-\alpha}\;p^{-\frac{1}{1-\alpha}-1},$$

which is strictly positive if and only if $p < \psi/\alpha = \alpha$. Operating profit is therefore strictly increasing in p on $(0,\alpha]$ and strictly decreasing for $p > \alpha$. If $G_{k,t}\,\psi < \alpha$, the patent holder posts the highest price at which farmers buy the improved vintage, $G_{k,t}\,\psi$. If $G_{k,t}\,\psi \geq \alpha$, it posts the unconstrained monopoly price α. Both the condition (A.3) and the price elasticity of demand are the same in the two countries, so the patent holder posts the same price in both.

The term $\Theta_{k,t}$ multiplies operating profit, so the optimal price does not depend on how many farmers chose type-k inputs or on where they are located. If no farmer chose type-k inputs, then $\Theta_{k,t} = 0$ and every price yields zero operating profit. In this case, I assume that the patent holder posts the same price, $\min\{\alpha, G_{k,t}\,\psi\}$. Finally, if type-$k$ innovation does not occur in period t, the competitive fringe supplies the most recent vintage at marginal cost ψ (Assumption A.2(a)). □

Two remarks are in order. First, the unconstrained monopoly price equals $\psi/\alpha = \alpha$ because farmers hire labor before prices are posted (Assumption A.2(b)), so the price elasticity of demand for intermediate inputs is $1/(1-\alpha)$. Second, in equilibrium, farmers anticipate the posted prices. Their choices in Steps 1 and 3 therefore coincide with the solution to the farmer's problem in Section 2.1.1, evaluated at the posted prices.

Margin rate and operating profit. By the demand for intermediate inputs, farmers in country j buy $x_{j,k,t} = \alpha\,\omega_{j,k,t}\,R_{j,t}/p^x_{k,t}$ units of type-k inputs. The operating profit that the patent holder earns in country j is therefore

$$\left(p^x_{k,t}-\psi\right)x_{j,k,t} = \frac{p^x_{k,t}-\psi}{p^x_{k,t}}\cdot\alpha\,\omega_{j,k,t}\,R_{j,t}.$$

By Lemma A.1, the first factor equals $1-1/G_{k,t}$ if innovation is non-drastic and $1-\psi/\alpha = 1-\alpha$ if it is drastic, which is the margin rate $\nu(G_{k,t})$ in (9). Summing over the two countries gives the profit function (8).

Proof of Lemma 1. Suppose that type-S innovation occurs in period t. Dividing the first line of (10) by $q_{S,t-1}$ and using $\lambda_{t-1} = q_{S,t-1}/q_{N,t-1}$ gives

$$G_{S,t} = \frac{q_{S,t}}{q_{S,t-1}} = (1+\xi)\left[(1-a)+\frac{a\chi}{\lambda_{t-1}}\right] = (1+\xi)\,\frac{T(\lambda_{t-1})}{\lambda_{t-1}},$$

which is (18). Because $a > 0$, G_S is strictly decreasing in λ_{t-1}. Moreover, $T(\chi) = \chi$ implies $G_S(\chi) = 1 + \xi$.

By definition, type-S innovation is drastic if and only if $G_S(\lambda_{t-1}) \geq 1/\alpha$ (Section 2.1.2), that is, if and only if

$$\alpha(1+\xi)\left[a\chi + (1-a)\lambda_{t-1}\right] \geq \lambda_{t-1} \qquad \Longleftrightarrow \qquad \lambda_{t-1} \leq \frac{\alpha(1+\xi)\,a\chi}{1-\alpha(1+\xi)(1-a)} = \lambda^{\mathrm{D}},$$

where the denominator is strictly positive because $\alpha(1+\xi) < 1$. The same inequality implies $\lambda^{\mathrm{D}} < \chi$.

The margin rate follows from (9). If $\lambda_{t-1} \leq \lambda^{\mathrm{D}}$, innovation is drastic and $\nu_{S,t} = 1 - \alpha$. If $\lambda^{\mathrm{D}} < \lambda_{t-1} \leq \chi$, innovation is non-drastic and $\nu_{S,t} = 1 - 1/G_S(\lambda_{t-1}) = 1 - \lambda_{t-1}/[(1+\xi)\,T(\lambda_{t-1})]$. The two expressions coincide at $\lambda_{t-1} = \lambda^{\mathrm{D}}$, where $G_S(\lambda^{\mathrm{D}}) = 1/\alpha$. Because G_S is strictly decreasing, the margin rate is strictly decreasing in λ_{t-1} on $[\lambda^{\mathrm{D}}, \chi]$. At $\lambda_{t-1} = \chi$, it equals $1 - 1/(1+\xi) = \xi/(1+\xi)$, which is the type-$N$ margin rate $\nu_{N,t} = \nu(1+\xi)$. □

A.3 Characterization of the Equilibrium in Period t

This section characterizes the equilibrium of period t, given the inherited quality ratio λ_{t-1} and the innovation decisions of the period. Under Assumption A.3(b), type-N innovation occurs in every period, so $m_{N,t} = 1$. The equilibrium of period t therefore depends on λ_{t-1} and on whether type-S innovation occurs, $m_{S,t} \in \{0,1\}$.

Proposition A.1. *Suppose that Assumptions A.1–A.4 hold. Fix a period t, the inherited quality ratio $\lambda_{t-1} \in (0,\chi]$, and $m_{S,t} \in \{0,1\}$. The equilibrium of period t satisfies the following.*

(i) Price-adjusted quality ratio. *The price-adjusted quality of type-N inputs is $q_{N,t}/p^x_{N,t} = q_{N,t-1}/\psi$. The price-adjusted quality of type-S inputs is $q_{S,t}/p^x_{S,t} = q_{S,t-1}/\psi$ if type-S innovation does not occur or is non-drastic, and $q_{S,t}/p^x_{S,t} = q_{S,t}/\alpha$ if it is drastic. Hence*

$$\widehat{\lambda}_t = \begin{cases} \lambda_{t-1} & \text{if } m_{S,t} = 0, \text{ or } m_{S,t} = 1 \text{ and } \lambda_{t-1} > \lambda^{D}, \\ \alpha(1+\xi)\,T(\lambda_{t-1}) & \text{if } m_{S,t} = 1 \text{ and } \lambda_{t-1} \leq \lambda^{D}, \end{cases}$$

which is (21). In both cases, $\lambda_{t-1} \leq \widehat{\lambda}_t \leq \chi$, and the two expressions coincide at $\lambda_{t-1} = \lambda^{D}$. If $a < 1$, each expression is strictly increasing in λ_{t-1}. If $a = 1$, the second expression is constant and equal to λ^{D}.

(ii) Adoption. *Northern farmers use type-N inputs: $\omega_{N,N,t} = 1$ and $\omega_{N,S,t} = 0$. Southern farmers use type-S inputs if and only if $\widehat{\lambda}_t \geq e^{-\delta}$: $\omega_{S,S,t} = \mathbb{1}\{\widehat{\lambda}_t \geq e^{-\delta}\}$ and $\omega_{S,N,t} = 1 - \omega_{S,S,t}$. If Assumption A.5(a) also holds, then $\widehat{\lambda}_t \geq e^{-\delta}$ if and only if $\lambda_{t-1} \geq e^{-\delta}$, whether or not type-S innovation occurs in period t.*

(iii) Price-adjusted effective quality of the inputs in use. $h^*_{N,t} = q_{N,t-1}/\psi$ *and* $h^*_{S,t} = \mathcal{Q}(\widehat{\lambda}_t)\, h^*_{N,t}$, *where* $\mathcal{Q}$ *is the function defined in Proposition 1.*

(iv) Relative price and agricultural revenue. *Let*

$$\mathcal{F}_t := \left(\frac{A^a_S}{A^a_N}\right)\left(\frac{A^n_S}{A^n_N}\right)^{-\beta} \mathcal{Q}(\widehat{\lambda}_t)^{\alpha}. \tag{A.4}$$

Then

$$\frac{p^a_{S,t}}{p^a_{N,t}} = \mathcal{F}_t^{-1/\Gamma}, \qquad \frac{R_{S,t}}{R_{N,t}} = \mathcal{F}_t^{(\varepsilon-1)/\Gamma}, \qquad \text{and} \qquad s_{S,t} = \left[1 + \mathcal{F}_t^{-(\varepsilon-1)/\Gamma}\right]^{-1}.$$

The share $s_{S,t}$ *is continuous in* $\widehat{\lambda}_t$. *It is constant in* $\widehat{\lambda}_t$ *for* $\widehat{\lambda}_t < e^{-\delta}$ *and strictly increasing in* $\widehat{\lambda}_t$ *for* $\widehat{\lambda}_t \geq e^{-\delta}$.

(v) World agricultural revenue.

$$R^W_t = \frac{\phi\left[A^n_N L_N + A^n_S L_S - c\,(1 + m_{S,t})\right]}{\kappa_t}, \tag{A.5}$$

where

$$\kappa_t := 1 - \phi\gamma - \phi\,\alpha\left[\nu_{N,t}\,(1 - \omega_{S,S,t}\, s_{S,t}) + m_{S,t}\,\nu_{S,t}\,\omega_{S,S,t}\, s_{S,t}\right] > 0. \tag{A.6}$$

(vi) Profits. *The operating profit from type-N innovation is* $\pi_{N,t} = \nu_{N,t}\,\alpha\left[R_{N,t} + (1 - \omega_{S,S,t})\, R_{S,t}\right]$. *If type-S innovation occurs, its operating profit is* $\pi_{S,t} = \nu_{S,t}\,\alpha\,\omega_{S,S,t}\, R_{S,t}$. *The net profit from type-k innovation is* $\Pi_{k,t} = \pi_{k,t} - c$. *If type-S innovation does not occur, no R&D firm earns profits from type-S inputs. In both cases, the profits of R&D firms in household income (16) are* $\sum_k m_{k,t}\,\Pi_{k,t}$.

Proof of part (i). Type-N innovation occurs in every period and is non-drastic (Assumption A.3(a)–(b)). By Lemma A.1, $p^x_{N,t} = (1+\xi)\,\psi$, and hence $q_{N,t}/p^x_{N,t} = (1+\xi)\,q_{N,t-1}/[(1+\xi)\,\psi] = q_{N,t-1}/\psi$.

For type-S inputs, there are three cases. First, if type-S innovation does not occur, the competitive fringe supplies the vintage of base quality $q_{S,t} = q_{S,t-1}$ at price ψ, so $q_{S,t}/p^x_{S,t} = q_{S,t-1}/\psi$. Second, if type-$S$ innovation occurs and is non-drastic, then $p^x_{S,t} = G_{S,t}\,\psi$ and $q_{S,t}/p^x_{S,t} = q_{S,t}/(G_{S,t}\,\psi) = q_{S,t-1}/\psi$. Third, if type-$S$ innovation occurs and is drastic, then $p^x_{S,t} = \alpha$ and $q_{S,t}/p^x_{S,t} = q_{S,t}/\alpha$. Dividing by $q_{N,t}/p^x_{N,t} = q_{N,t-1}/\psi$ gives $\widehat{\lambda}_t = q_{S,t-1}/q_{N,t-1} = \lambda_{t-1}$ in the first two cases. In the third case, the catch-up law (10) gives $q_{S,t} = (1+\xi)\,T(\lambda_{t-1})\,q_{N,t-1}$, and hence

$$\widehat{\lambda}_t = \frac{q_{S,t}}{\alpha}\cdot\frac{\psi}{q_{N,t-1}} = \frac{\alpha\, q_{S,t}}{q_{N,t-1}} = \alpha(1+\xi)\,T(\lambda_{t-1}).$$

It remains to show the properties of $\widehat{\lambda}_t$. By Lemma 1, type-S innovation is drastic if and only if $G_S(\lambda_{t-1}) \geq 1/\alpha$, that is, $\alpha(1+\xi)\,T(\lambda_{t-1}) \geq \lambda_{t-1}$, with equality if and only if $\lambda_{t-1} = \lambda^{\mathrm{D}}$. Hence $\widehat{\lambda}_t \geq \lambda_{t-1}$, and the two expressions coincide at $\lambda_{t-1} = \lambda^{\mathrm{D}}$. Moreover, $\lambda_{t-1} \leq \chi$ by hypothesis, and $\alpha(1+\xi)\,T(\lambda_{t-1}) \leq \alpha(1+\xi)\,\chi < \chi$ by Assumption A.3(a), so $\widehat{\lambda}_t \leq \chi$. Finally, $T(\lambda_{t-1}) = (1-a)\,\lambda_{t-1} + a\chi$ is strictly increasing in λ_{t-1} if $a < 1$. If $a = 1$, then $T(\lambda_{t-1}) = \chi$ for every λ_{t-1}, so the second expression equals $\alpha(1+\xi)\,\chi$, which is the value of λ^{D} at $a = 1$. □

Proof of part (ii). By the definitions of $h_{j,k,t}$ and $\widehat{\lambda}_t$, the Northern farmer compares $h_{N,N,t} = q_{N,t}/p^x_{N,t}$ with $h_{N,S,t} = e^{-\delta} q_{S,t}/p^x_{S,t} = e^{-\delta}\widehat{\lambda}_t h_{N,N,t}$. Because $\widehat{\lambda}_t \le \chi \le 1$ by part (i), $h_{N,S,t} < h_{N,N,t}$, so Northern farmers use type-N inputs. The Southern farmer compares $h_{S,N,t} = e^{-\delta} q_{N,t}/p^x_{N,t}$ with $h_{S,S,t} = q_{S,t}/p^x_{S,t} = e^{\delta}\widehat{\lambda}_t h_{S,N,t}$. Hence $h_{S,S,t} \ge h_{S,N,t}$ if and only if $\widehat{\lambda}_t \ge e^{-\delta}$, and if the two are equal the Southern farmer uses type-S inputs (Assumption A.1).

Suppose now that Assumption A.5(a) holds. If type-S innovation does not occur or is non-drastic, then $\widehat{\lambda}_t = \lambda_{t-1}$ by part (i), and the claim is immediate. The same is true if $\lambda_{t-1} = \lambda^{\mathrm{D}}$. Suppose finally that type-S innovation occurs and $\lambda_{t-1} < \lambda^{\mathrm{D}}$, so that $\lambda_{t-1} < e^{-\delta}$. If $a < 1$, then $\widehat{\lambda}_t = \alpha(1+\xi)\,T(\lambda_{t-1}) < \alpha(1+\xi)\,T(\lambda^{\mathrm{D}}) = \lambda^{\mathrm{D}} \le e^{-\delta}$. If $a = 1$, then $\widehat{\lambda}_t = \lambda^{\mathrm{D}} < e^{-\delta}$. In both cases, $\widehat{\lambda}_t < e^{-\delta}$ and $\lambda_{t-1} < e^{-\delta}$. □

Proof of part (iii). By part (ii), Northern farmers use type-N inputs, so $h^*_{N,t} = h_{N,N,t} = q_{N,t}/p^x_{N,t} = q_{N,t-1}/\psi$ by part (i). Southern farmers use the input type with the higher price-adjusted effective quality, so

$$h^*_{S,t} = \max\{h_{S,N,t},\, h_{S,S,t}\} = \frac{q_{N,t}}{p^x_{N,t}} \max\{e^{-\delta},\, \widehat{\lambda}_t\} = \mathcal{Q}(\widehat{\lambda}_t)\, h^*_{N,t},$$

because $\mathcal{Q}(\widehat{\lambda}_t) = e^{-\delta}$ if $\widehat{\lambda}_t < e^{-\delta}$ and $\mathcal{Q}(\widehat{\lambda}_t) = \widehat{\lambda}_t$ if $\widehat{\lambda}_t \ge e^{-\delta}$. □

Proof of part (iv). Demand. The household in country j chooses $c^a_{j,N,t}$ and $c^a_{j,S,t}$ to maximize $C^a_{j,t}$ in (14) for a given expenditure on agricultural goods. With $\mu = 1/2$, the first-order conditions give $c^a_{j,S,t}/c^a_{j,N,t} = (p^a_{S,t}/p^a_{N,t})^{-\varepsilon}$ for $j \in \{N,S\}$. Summing over the two countries and using the market clearing conditions for agricultural goods gives $Y^a_{S,t}/Y^a_{N,t} = (p^a_{S,t}/p^a_{N,t})^{-\varepsilon}$, and hence

$$\frac{R_{S,t}}{R_{N,t}} = \left(\frac{p^a_{S,t}}{p^a_{N,t}}\right)^{1-\varepsilon}.$$

Supply. Dividing (A.1) for the South by (A.1) for the North, and using $w^a_{j,t} = A^n_j$ and part (iii), gives

$$\left(\frac{R_{S,t}}{R_{N,t}}\right)^{\gamma} = \left(\frac{A^a_S}{A^a_N}\right)\left(\frac{A^n_S}{A^n_N}\right)^{-\beta}\left(\frac{h^*_{S,t}}{h^*_{N,t}}\right)^{\alpha}\frac{p^a_{S,t}}{p^a_{N,t}} = \mathcal{F}_t\,\frac{p^a_{S,t}}{p^a_{N,t}}.$$

Equilibrium. Combining the two equations gives $(p^a_{S,t}/p^a_{N,t})^{\gamma(1-\varepsilon)} = \mathcal{F}_t\,(p^a_{S,t}/p^a_{N,t})$, that is, $(p^a_{S,t}/p^a_{N,t})^{-\Gamma} = \mathcal{F}_t$. Hence $p^a_{S,t}/p^a_{N,t} = \mathcal{F}_t^{-1/\Gamma}$ and $R_{S,t}/R_{N,t} = \mathcal{F}_t^{(\varepsilon-1)/\Gamma}$, and the expression for $s_{S,t} = [1 + R_{N,t}/R_{S,t}]^{-1}$ follows.

Properties of $s_{S,t}$. The function $\mathcal{Q}$ is continuous, constant and equal to $e^{-\delta}$ for $\widehat{\lambda}_t < e^{-\delta}$, and strictly increasing for $\widehat{\lambda}_t \ge e^{-\delta}$. Because $\varepsilon > 1$, $s_{S,t}$ is strictly increasing in $\mathcal{F}_t$, and $\mathcal{F}_t$ is strictly increasing in $\mathcal{Q}(\widehat{\lambda}_t)$. The stated properties follow. □

Proof of parts (v) and (vi). I begin with part (vi). By part (ii), type-N inputs are used by Northern farmers and, if $\omega_{S,S,t} = 0$, by Southern farmers. By (8), $\pi_{N,t} = \nu_{N,t}\,\alpha\,[R_{N,t} + (1-\omega_{S,S,t})\,R_{S,t}]$, and, if type-$S$ innovation occurs, $\pi_{S,t} = \nu_{S,t}\,\alpha\,\omega_{S,S,t}\,R_{S,t}$. If type-$S$ innovation does not occur, the

competitive fringe supplies type-S inputs at marginal cost, no R&D firm enters, and no R&D fixed cost is paid for type-S inputs. The profits of R&D firms received by Northern households are therefore $\sum_k m_{k,t}\,\Pi_{k,t}$, with $\Pi_{k,t} = \pi_{k,t} - c$.

For part (v), consider household income (16). With equalized wages, labor income is $A_j^n L_j$ in country j, which in the North includes the income of researchers. The profits of agricultural final good producers are $\gamma R_{j,t}$. Summing over the two countries and using $R_{N,t} + (1-\omega_{S,S,t})\,R_{S,t} = (1-\omega_{S,S,t}\,s_{S,t})\,R_t^W$ and $\omega_{S,S,t}\,R_{S,t} = \omega_{S,S,t}\,s_{S,t}\,R_t^W$ gives

$$E_{N,t} + E_{S,t} = A_N^n L_N + A_S^n L_S + \gamma R_t^W + \alpha\left[\nu_{N,t}\,(1-\omega_{S,S,t}\,s_{S,t}) + m_{S,t}\,\nu_{S,t}\,\omega_{S,S,t}\,s_{S,t}\right] R_t^W - c\,(1+m_{S,t}).$$

With the preferences in (14), each household spends the share ϕ of its income on agricultural goods. The market clearing conditions for agricultural goods then imply $R_t^W = \phi\,(E_{N,t} + E_{S,t})$. Substituting the expression for $E_{N,t} + E_{S,t}$ and solving for R_t^W gives (A.5).

It remains to show that $\kappa_t > 0$. Every margin rate is at most $1-\alpha$ by (9), and the two weights in the square brackets of (A.6) are nonnegative and sum to at most one. Hence $\kappa_t \geq 1 - \phi\gamma - \phi\alpha(1-\alpha) > 1 - \phi\,(\gamma+\alpha) > 0$. The numerator of (A.5) is strictly positive by Assumption A.4(b). □

Three remarks are in order. First, $s_{S,t}$ and κ_t depend on the inherited quality ratio only through $\widehat{\lambda}_t$ and the margin rate $\nu_{S,t}$, and not on the level of world agricultural revenue. Equation (A.5) therefore determines R_t^W in closed form. Second, if $\widehat{\lambda}_t < e^{-\delta}$, farmers in both countries use type-N inputs, and the relative price of agricultural goods, the share $s_{S,t}$, and κ_t do not depend on λ_{t-1}. The quality ratio may change from one period to the next, but these objects, and therefore agricultural productivity differences (22), remain constant as long as the price-adjusted quality ratio remains below $e^{-\delta}$. Third, the case $a = 1$ differs from the case $a < 1$ in one respect. If $a = 1$, type-S innovation raises the quality ratio to χ regardless of the inherited quality ratio, so the price-adjusted quality ratio after drastic innovation is the same at every $\lambda_{t-1} \leq \lambda^{\mathrm{D}}$ and equals λ^{D}. This is why Assumption A.5(a) requires a strict inequality if $a = 1$: if $\lambda^{\mathrm{D}} = e^{-\delta}$ and $a = 1$, Southern farmers would adopt type-S inputs after drastic innovation by Assumption A.1.

The following corollary expresses the profit from type-S innovation as a function of the inherited quality ratio, as in Section 2.3.2.

Corollary A.1. *Suppose that Assumptions A.1–A.4 hold. Let $\pi_{S,t}(\lambda_{t-1})$ denote the operating profit of an R&D firm that innovates type-S inputs in period t, given the inherited quality ratio λ_{t-1}, and let $\Pi_{S,t}(\lambda_{t-1}) = \pi_{S,t}(\lambda_{t-1}) - c$. Then*

$$\pi_{S,t}(\lambda_{t-1}) = \nu\big(G_S(\lambda_{t-1})\big) \cdot \alpha\,\mathbb{1}\{\widehat{\lambda}_t \geq e^{-\delta}\}\,s_{S,t}\,R_t^W,$$

where $\widehat{\lambda}_t$, $s_{S,t}$, and R_t^W are given by Proposition A.1 with $m_{S,t} = 1$. If Assumption A.5(a) also holds, then

$$\pi_{S,t}(\lambda_{t-1}) = \begin{cases} \nu\big(G_S(\lambda_{t-1})\big) \cdot \alpha\,s_{S,t}\,R_t^W & \text{if } \lambda_{t-1} \geq e^{-\delta}, \\ 0 & \text{if } \lambda_{t-1} < e^{-\delta}, \end{cases} \tag{A.7}$$

where $s_{S,t}$ and R_t^W are evaluated at $\widehat{\lambda}_t = \lambda_{t-1}$. In particular, $\Pi_{S,t}(\lambda_{t-1}) = -c$ for every $\lambda_{t-1} < e^{-\delta}$.

Proof. The first expression follows from Proposition A.1(ii) and (vi) and the definition of $s_{S,t}$. Under Assumption A.5(a), $\widehat{\lambda}_t \geq e^{-\delta}$ if and only if $\lambda_{t-1} \geq e^{-\delta}$ by Proposition A.1(ii). If $\lambda_{t-1} \geq e^{-\delta}$, then $\lambda_{t-1} \geq \lambda^{\mathrm{D}}$, so $\widehat{\lambda}_t = \lambda_{t-1}$ by Proposition A.1(i). □

The function $\pi_{S,t}(\lambda_{t-1})$ is the operating profit that an R&D firm anticipates when it decides whether to innovate type-S inputs. It differs from the operating profit earned in equilibrium, which is zero in every period in which type-S innovation does not occur. If $\widehat{\lambda}_t < e^{-\delta}$, type-$S$ inputs have no market even after they are improved, because Southern farmers do not adopt them. If innovation is non-drastic, the reason is that the patent holder charges exactly for the quality improvement: at the price $G_{S,t}\psi$, the improved vintage has the same price-adjusted quality as the previous vintage at price ψ, which Southern farmers do not use.

A.4 Uniqueness of the Equilibrium

Lemma A.2. *Suppose that Assumptions A.1–A.4 hold. Fix a period t, the inherited quality ratio $\lambda_{t-1} \in (0,\chi]$, and $m_{S,t} \in \{0,1\}$. The equilibrium of period t exists and is unique. Moreover, the equilibrium path from any initial quality ratio $\lambda_{-1} \in (0,\chi]$ is unique.*

Proof. By Lemma A.1, the prices of intermediate inputs are uniquely determined by the quality improvements $G_{N,t}$ and $G_{S,t}$, and they do not depend on which input types farmers choose. Farmers anticipate these prices in Step 1 of Stage 2. Given the prices, the choice of the farmer in country j depends only on the comparison of $h_{j,N,t}$ with $h_{j,S,t}$, which does not involve the choices of other farmers, the agricultural output price, or the wage. The input type is therefore uniquely determined if $h_{j,N,t} \neq h_{j,S,t}$, and by Assumption A.1 otherwise. Given the input types, Proposition A.1(iii)–(v) determines the relative price of agricultural goods, the share $s_{S,t}$, and world agricultural revenue R_t^W uniquely. The revenues $R_{j,t}$, the agricultural output prices from (A.1), agricultural output and employment, household incomes, and consumption then follow. Hence the equilibrium of period t exists and is unique. In particular, there is no equilibrium in which Southern farmers do not use type-S inputs only because other Southern farmers do not use them.

For the equilibrium path, fix λ_{t-1}. The net profit $\Pi_{S,t}(\lambda_{t-1})$ in Corollary A.1 is uniquely determined, and type-S innovation occurs if and only if $\Pi_{S,t}(\lambda_{t-1}) \geq 0$ by the entry condition (12). The law of motion (17) then determines λ_t uniquely. Starting from λ_{-1}, the path of the quality ratio and the equilibrium of every period are therefore unique. □

A.5 Proof of Proposition 1 and Agricultural Employment Shares

Proof of Proposition 1. By (A.2) and Proposition A.1(iv),

$$\frac{Y^a_{S,t}/L^a_{S,t}}{Y^a_{N,t}/L^a_{N,t}} = \frac{A^n_S}{A^n_N}\cdot\frac{p^a_{N,t}}{p^a_{S,t}} = \frac{A^n_S}{A^n_N}\mathcal{F}_t^{1/\Gamma} = \left[\left(\frac{A^n_S}{A^n_N}\right)^{\Gamma}\left(\frac{A^a_S}{A^a_N}\right)\left(\frac{A^n_S}{A^n_N}\right)^{-\beta}\mathcal{Q}(\widehat{\lambda}_t)^{\alpha}\right]^{1/\Gamma}.$$

Because $\Gamma-\beta = 1+\gamma\varepsilon-\gamma-\beta = \alpha+\gamma\varepsilon$, this is (22). □

The proof uses Assumptions A.1–A.4 only through (A.2) and Proposition A.1(ii)–(iv). In particular, (22) holds whether type-S innovation is drastic or non-drastic, because the pricing of intermediate inputs affects agricultural productivity differences only through $\widehat{\lambda}_t$.

The same steps give the ratio of agricultural employment shares between the South and the North.

Corollary A.2. *Suppose that Assumptions A.1–A.4 hold. In equilibrium, the ratio of agricultural employment shares between the South and the North in period t is given by*

$$\frac{L^a_{S,t}/L_S}{L^a_{N,t}/L_N} = \frac{L_N}{L_S}\left(\frac{A^n_S}{A^n_N}\right)^{-\left[1+\frac{\beta(\varepsilon-1)}{\Gamma}\right]}\left(\frac{A^a_S}{A^a_N}\right)^{\frac{\varepsilon-1}{\Gamma}}\mathcal{Q}(\widehat{\lambda}_t)^{\frac{\alpha(\varepsilon-1)}{\Gamma}}. \tag{A.8}$$

Proof. By (A.2), $L^a_{j,t}/L_j = \beta R_{j,t}/(A^n_j L_j)$. By Proposition A.1(iv),

$$\frac{L^a_{S,t}/L_S}{L^a_{N,t}/L_N} = \frac{L_N}{L_S}\left(\frac{A^n_S}{A^n_N}\right)^{-1}\frac{R_{S,t}}{R_{N,t}} = \frac{L_N}{L_S}\left(\frac{A^n_S}{A^n_N}\right)^{-1}\mathcal{F}_t^{(\varepsilon-1)/\Gamma}.$$

Substituting (A.4) gives (A.8). □

Equation (A.8) shows how the three sources in Proposition 1 affect agricultural employment. Higher non-agricultural productivity in the South lowers the Southern agricultural employment share relative to the Northern one, which is the labor pull channel. Stronger agricultural fundamentals in the South and a higher ratio $\mathcal{Q}(\widehat{\lambda}_t)$ raise it. The reason is that the demand for each country's agricultural good is elastic ($\varepsilon > 1$): the country that produces agricultural goods more efficiently earns a larger share of world agricultural revenue and employs more agricultural labor. Environmental mismatch, which lowers $\mathcal{Q}(\widehat{\lambda}_t)$, therefore widens agricultural productivity differences and, at the same time, lowers agricultural employment in the South relative to the North.

A.6 Endogenous Location of R&D

Section 2 assumes that all R&D firms are located in the North (Assumption A.3(c)). This section allows R&D firms to be located in either country and shows that, in equilibrium, R&D takes place only in the North.

Consider the following extension of the baseline model. An R&D firm located in country j requires $1/\eta_j$ researchers to innovate either input type, where $\eta_j > 0$ denotes research productivity in country j. Its fixed cost of R&D is

$$c_j := \frac{w^r_{j,t}}{\eta_j}, \qquad j \in \{N, S\}, \tag{A.9}$$

where $w^r_{j,t}$ denotes the wage paid to researchers in country j. With equalized wages, $w^r_{j,t} = A^n_j$, so $c_j = A^n_j/\eta_j$, and c_N is the fixed cost c in (11) with $\eta_N = \eta$. In Stage 1, potential R&D firms in both countries decide whether to enter. As in Section 2.1.2, entry is sequential and observable, and the first firm to enter and improve type-k inputs receives the one-period patent. I assume that firms with a lower fixed cost decide first. Everything else is as in Section 2. In particular, the quality improvement, the price, and the market size of type-k inputs do not depend on the country in which the patent holder is located. The operating profit from type-k innovation is therefore $\pi_{k,t}$, regardless of which firm innovates.

Lemma A.3. *Suppose that the fixed cost of R&D is higher in the South than in the North:*

$$c_S > c_N \qquad \Longleftrightarrow \qquad \frac{\eta_S}{\eta_N} < \frac{A^n_S}{A^n_N}. \tag{A.10}$$

Then, in equilibrium, every innovation is undertaken by an R&D firm located in the North, and type-k innovation occurs in period t if and only if $\pi_{k,t} \geq c_N$.

Proof. Fix an input type k and a period t. If $\pi_{k,t} \geq c_N$, the Northern firm, which decides first, enters, because its net profit $\pi_{k,t} - c_N$ is nonnegative. The Southern firm then does not enter, because it would not receive the patent and could not recover its fixed cost of R&D. If $\pi_{k,t} < c_N$, the Northern firm does not enter. The Southern firm does not enter either, because $\pi_{k,t} < c_N < c_S$ implies that its net profit would be strictly negative. In both cases, no Southern firm enters, and type-k innovation occurs if and only if $\pi_{k,t} \geq c_N$, which is the entry condition (12). The innovation decisions, the base qualities, and the equilibrium of every period are therefore those of the baseline model. □

Condition (A.10) requires that the South's disadvantage in research productivity exceed its wage advantage: relative research productivity, η_S/η_N, must be lower than the relative wage, A^n_S/A^n_N. The assumption on the order of entry is not needed under a stronger condition. If c_S exceeds the operating profit from either input type in every period, no Southern firm enters, regardless of the order of entry. Finally, the location of R&D matters only for which country's households supply the researchers and receive the profits of R&D firms. If a Southern firm innovated type-S inputs at the same fixed cost, the quality ratio and agricultural productivity differences would be unchanged.

A.7 The Laissez-Faire Path

This section characterizes the path of the quality ratio under laissez-faire, which Section 2.3.2 describes in words: the set of quality ratios at which type-S innovation is profitable, the long-run outcomes, and the dynamic amplification factor. Throughout this section and the next three, Assumptions A.1–A.5 hold. I write $\pi_{S,t}(\lambda)$ for the operating profit of Corollary A.1 with $s_{S,t}$ and R_t^W evaluated at $\widehat{\lambda}_t = \lambda$ and $m_{S,t} = 1$, and $\Pi_{S,t}(\lambda) = \pi_{S,t}(\lambda) - c$.

Where type-S innovation is profitable. By Corollary A.1, $\pi_{S,t}(\lambda) = 0$ for $\lambda < e^{-\delta}$, so type-S innovation never occurs there: the improved inputs would have no buyer, and the R&D firm would pay the fixed cost for nothing. For $\lambda \in [e^{-\delta}, \chi]$, Southern farmers use type-S inputs and

$$\begin{aligned} \pi_{S,t}(\lambda) &= \nu\big(G_S(\lambda)\big) \cdot \alpha\, s_{S,t}(\lambda)\, R_t^W(\lambda) > 0, \\ s_{S,t}(\lambda) &= \Big[1 + \mathcal{F}(\lambda)^{-(\varepsilon-1)/\Gamma}\Big]^{-1}, \qquad R_t^W(\lambda) = \frac{\phi\left[A_N^n L_N + A_S^n L_S - 2c\right]}{\kappa(\lambda)}, \end{aligned} \tag{A.11}$$

where $\mathcal{F}(\lambda) = (A_S^a/A_N^a)(A_S^n/A_N^n)^{-\beta}\lambda^{\alpha}$ and $\kappa(\lambda) = 1 - \phi\gamma - \phi\alpha\left[\nu_N(1 - s_{S,t}(\lambda)) + \nu(G_S(\lambda))\, s_{S,t}(\lambda)\right]$ are the objects of Proposition A.1(iv)–(v) at $\widehat{\lambda}_t = \lambda$ and $\omega_{S,S,t} = 1$. The profit is continuous on $[e^{-\delta}, \chi]$, and the following lemma decomposes its slope into the two forces that Section 2.3.2 names: the margin rate, which falls as the quality ratio rises because the catch-up $a(\chi - \lambda)$ shrinks, and the market size, which grows with the quality ratio because the South's share of world agricultural revenue rises.

Lemma A.4. *On $(e^{-\delta}, \chi]$, $\pi_{S,t}$ is continuously differentiable, and*

$$\frac{d\ln \pi_{S,t}(\lambda)}{d\ln \lambda} = \mathcal{E}_\nu(\lambda) + \mathcal{E}_s(\lambda) + \mathcal{E}_R(\lambda), \tag{A.12}$$

where

$$\begin{aligned} \mathcal{E}_\nu(\lambda) &:= -\frac{a\chi}{T(\lambda)\left(G_S(\lambda) - 1\right)}, \qquad \mathcal{E}_s(\lambda) := \psi_e\left(1 - s_{S,t}(\lambda)\right), \\ \mathcal{E}_R(\lambda) &:= \frac{\phi\alpha\left[\left(\nu(G_S(\lambda)) - \nu_N\right)\psi_e\, s_{S,t}(\lambda)\left(1 - s_{S,t}(\lambda)\right) + s_{S,t}(\lambda)\,\nu(G_S(\lambda))\,\mathcal{E}_\nu(\lambda)\right]}{\kappa(\lambda)}, \end{aligned}$$

and $\psi_e := \alpha(\varepsilon - 1)/\Gamma$. The first term, the response of the margin rate, is negative; the second, the response of the South's share of world agricultural revenue, is positive and smaller than ψ_e; the third collects the response of world agricultural revenue. None of the three depends on c.

Proof. Write $\ln \pi_{S,t} = \ln \nu(G_S) + \ln s_{S,t} + \ln R_t^W$ and differentiate with respect to $\ln \lambda$. On $(e^{-\delta}, \chi]$ the innovation is non-drastic, since $\lambda > e^{-\delta} \geq \lambda^{\mathrm{D}}$, so $\nu(G_S) = 1 - 1/G_S$ and $d\ln\nu/d\ln\lambda = [1/(G_S - 1)]\, d\ln G_S/d\ln\lambda$; by (18), $d\ln G_S/d\ln\lambda = \lambda T'(\lambda)/T(\lambda) - 1 = -a\chi/T(\lambda)$, which gives the first term, negative because $G_S > 1$. For the share, $d\ln\mathcal{F}/d\ln\lambda = \alpha$, so with $X := \mathcal{F}^{-(\varepsilon-1)/\Gamma}$, $d\ln s_{S,t}/d\ln\lambda = -(1 - s_{S,t})\, d\ln X/d\ln\lambda = \psi_e(1 - s_{S,t}) \in (0, \psi_e)$. For world

revenue, $d\ln R_t^W/d\ln\lambda = -d\ln\kappa/d\ln\lambda$, and $d\kappa/d\ln\lambda = -\phi\alpha[(\nu(G_S)-\nu_N)\,ds_{S,t}/d\ln\lambda + s_{S,t}\,d\nu(G_S)/d\ln\lambda]$, which gives the third term. The cost enters R_t^W only through the level $A_N^n L_N + A_S^n L_S - 2c$, not through κ. □

The sum of the three terms is not signed by the primitives. The following proposition states the long-run outcomes when the profit falls as the quality ratio rises on $[e^{-\delta},\chi]$, so that the largest profit is earned at $e^{-\delta}$ and the smallest at χ; a remark after the proof records what changes when it does not.

Proposition A.2. *Suppose that $\pi_{S,t}$ is strictly decreasing on $[e^{-\delta},\chi]$ and that $c \notin \{\pi_{S,t}(e^{-\delta}),\pi_{S,t}(\chi)\}$. Along the unique laissez-faire path from any $\lambda_{-1} \in (0,\chi]$ (Lemma A.2), type-S innovation never occurs in a period in which $\lambda_{t-1} < e^{-\delta}$, and the quality ratio then declines as $\lambda_t = \lambda_{t-1}/(1+\xi)$. With type-S innovation in every period, $\lambda_t = \chi - (1-a)^{t+1}(\chi-\lambda_{-1})$, which rises to χ; without it, $\lambda_t = \lambda_{-1}/(1+\xi)^{t+1}$, which declines to zero. Exactly one of the following holds.*

- *(a) If $c < \pi_{S,t}(\chi)$, type-S innovation is profitable at every $\lambda \in [e^{-\delta},\chi]$. From $\lambda_{-1} < e^{-\delta}$ it never occurs, Southern farmers use type-N inputs in every period, and the quality ratio declines to zero. From $\lambda_{-1} \geq e^{-\delta}$ it occurs in every period, and the quality ratio rises to χ.*

- *(b) If $\pi_{S,t}(\chi) < c < \pi_{S,t}(e^{-\delta})$, type-S innovation is profitable exactly on $[e^{-\delta},\bar\lambda]$, where $\bar\lambda \in (e^{-\delta},\chi)$ is the unique solution of $\pi_{S,t}(\bar\lambda) = c$. From $\lambda_{-1} < e^{-\delta}$ the path is as in (a). From $\lambda_{-1} \geq e^{-\delta}$, if $\bar\lambda \geq (1+\xi)\,e^{-\delta}$, type-S innovation occurs in infinitely many periods but not in every period: innovation carries the quality ratio above $\bar\lambda$, where the profit no longer covers the fixed cost, the ratio then declines back to $[e^{-\delta},\bar\lambda]$, and innovation resumes.*

- *(c) If $c > \pi_{S,t}(e^{-\delta})$, type-S innovation never occurs, from any λ_{-1}, and the quality ratio declines to zero.*

Proof. The closed forms. T is affine with slope $1-a$ and fixed point χ, so $T^n(\lambda)-\chi = (1-a)^n(\lambda-\chi)$ by induction, which with $\lambda_t = T^{t+1}(\lambda_{-1})$ gives the first form; the second is the decay branch of (17) iterated. *The set of profitable quality ratios.* Let $\mathcal{S} := \{\lambda \in (0,\chi] : \pi_{S,t}(\lambda) \geq c\}$. By Corollary A.1, $\mathcal{S} \cap (0,e^{-\delta}) = \varnothing$. On $[e^{-\delta},\chi]$, $\pi_{S,t}$ is continuous and strictly decreasing from $\pi_{S,t}(e^{-\delta})$ to $\pi_{S,t}(\chi)$, so with c different from both values: $\mathcal{S} = [e^{-\delta},\chi]$ in case (a); $\mathcal{S} = [e^{-\delta},\bar\lambda]$ in case (b), with $\bar\lambda$ the unique solution of $\pi_{S,t}(\bar\lambda) = c$; and $\mathcal{S} = \varnothing$ in case (c). The quality ratio $e^{-\delta}$ belongs to $\mathcal{S}$ in the first two cases because Southern farmers adopt type-S inputs there (Assumption A.1) and an indifferent firm enters. *The dynamics.* By (12), type-S innovation occurs in period t if and only if $\lambda_{t-1} \in \mathcal{S}$. From any $\lambda_{t-1} < e^{-\delta}$ there is no innovation, the ratio declines, and it never re-enters $\mathcal{S}$, which proves the statements about starts below $e^{-\delta}$ and all of case (c). In case (a), from $\lambda_{t-1} \geq e^{-\delta}$ innovation occurs and $T(\lambda_{t-1}) \in [e^{-\delta},\chi]$, since $T(\lambda) \geq \lambda$ and $T(\chi) = \chi$; so T maps $\mathcal{S}$ into itself, innovation occurs in every period, and the ratio rises to χ. In case (b), from $\lambda_{t-1} \in [e^{-\delta},\bar\lambda]$ each innovation raises the ratio by $T(\lambda)-\lambda = a(\chi-\lambda) \geq a(\chi-\bar\lambda) > 0$, so finitely many innovations carry it above $\bar\lambda$, where innovation stops and the

ratio declines. If $\bar{\lambda} \geq (1+\xi)\,e^{-\delta}$, a decline that starts above $\bar{\lambda}$ lands at or above $\bar{\lambda}/(1+\xi) \geq e^{-\delta}$, so the first landing at or below $\bar{\lambda}$ lies in $\mathcal{S}$ and innovation resumes; the interval $[\bar{\lambda}/(1+\xi), T(\bar{\lambda})]$ is mapped into itself by the laissez-faire path and absorbs every path from $\lambda_{-1} \geq e^{-\delta}$ in finite time. Innovation occurs infinitely often, because every stay above $\bar{\lambda}$ ends in a landing inside $\mathcal{S}$, and never in every period, because $\Pi_{S,t}(\chi) < 0$ and χ is the only fixed point of T. The three cases partition the position of c relative to the two values, so exactly one holds. □

Two remarks are in order. First, when $\bar{\lambda} < (1+\xi)\,e^{-\delta}$ in case (b), a decline from just above $\bar{\lambda}$ can land below $e^{-\delta}$, and the long run from $\lambda_{-1} \geq e^{-\delta}$ is not determined by the position of c alone: depending on $\bar{\lambda}$, the path executes finitely many bursts of innovation and then declines to zero, or sustains repeated bursts. Second, if $\pi_{S,t}$ rises on part of $[e^{-\delta}, \chi]$, the set $\mathcal{S}$ can be an interval whose lower end lies above $e^{-\delta}$, at quality ratios where Southern farmers use type-S inputs but the profit on the whole Southern market falls short of the fixed cost; the statements about starts below $e^{-\delta}$ are unchanged, since they rest on Corollary A.1 alone.

The number of innovations that reach $e^{-\delta}$. For $\lambda < e^{-\delta}$, let

$$T^*(\lambda) := \begin{cases} \left\lceil \dfrac{\ln\left[(\chi - e^{-\delta})/(\chi - \lambda)\right]}{\ln(1-a)} \right\rceil & \text{if } a < 1, \\ 1 & \text{if } a = 1. \end{cases} \tag{A.13}$$

By the closed form of Proposition A.2, $T^n(\lambda) = \chi - (1-a)^n(\chi - \lambda) \geq e^{-\delta}$ if and only if $(1-a)^n \leq (\chi - e^{-\delta})/(\chi - \lambda)$, which for $a < 1$ holds if and only if $n \geq \ln[(\chi - e^{-\delta})/(\chi - \lambda)]/\ln(1-a)$, since $\ln(1-a) < 0$; for $a = 1$, $T(\lambda) = \chi > e^{-\delta}$. Hence $T^*(\lambda)$ is the number of consecutive type-S innovations that lift the quality ratio from λ to $e^{-\delta}$ or above: along the path on which type-S inputs are improved in every period from $\lambda_{-1} < e^{-\delta}$, the inherited quality ratio $T^t(\lambda_{-1})$ lies below $e^{-\delta}$ exactly for $t < T^*(\lambda_{-1})$, so Southern farmers use type-N inputs in those periods and type-S inputs from period $T^*(\lambda_{-1})$ on.

The dynamic amplification factor. Section 2.3.2 defines $\mathcal{A}_t = \mathrm{APD}_t^U / \mathrm{APD}_t^D = [\mathcal{Q}(\widehat{\lambda}_t^U)/\mathcal{Q}(\widehat{\lambda}_t^D)]^{\alpha/\Gamma}$, where D is the laissez-faire path and U the path on which both input types are improved in every period, from a common λ_{-1}, and states that along the lock-in path the factor is nondecreasing and converges to $(\chi/e^{-\delta})^{\alpha/\Gamma}$. The following proposition proves these statements, with $T^*(\lambda)$ as in (A.13).

Proposition A.3. *Suppose that type-S innovation never occurs along the laissez-faire path from λ_{-1}, which is the case for every $\lambda_{-1} < e^{-\delta}$ and, in case (c) of Proposition A.2, for every $\lambda_{-1} \in (0, \chi]$. Then, for every $t \geq 0$, $\mathcal{Q}(\widehat{\lambda}_t^k) = \mathcal{Q}(\lambda_{t-1}^k)$ on both paths $k \in \{D, U\}$, and*

$$\mathcal{A}_t = \left[\frac{\mathcal{Q}\big(T^t(\lambda_{-1})\big)}{\mathcal{Q}\big(\lambda_{-1}/(1+\xi)^t\big)} \right]^{\alpha/\Gamma}$$

is nondecreasing in t and converges to $(\chi/e^{-\delta})^{\alpha/\Gamma}$. If $\lambda_{-1} < e^{-\delta}$, then $\mathcal{A}_t = 1$ for $t < T^*(\lambda_{-1})$ and $\mathcal{A}_t = [T^t(\lambda_{-1})/e^{-\delta}]^{\alpha/\Gamma} \geq 1$ from $t = T^*(\lambda_{-1})$ on, with strict inequality for $t > T^*(\lambda_{-1})$. If instead type-S innovation occurs in every period along the laissez-faire path from λ_{-1}, the two paths coincide and $\mathcal{A}_t = 1$ for every t.

Proof. The price-adjusted and the inherited ratio. If type-S innovation does not occur or is non-drastic, $\widehat{\lambda}_t = \lambda_{t-1}$ by Proposition A.1(i). If it is drastic, then $\lambda_{t-1} \leq \lambda^{\mathrm{D}} \leq e^{-\delta}$ and $\lambda_{t-1} \leq \widehat{\lambda}_t \leq \lambda^{\mathrm{D}}$ (Proposition A.1(i)), so both ratios lie in $(0, e^{-\delta}]$, where $\mathcal{Q} = e^{-\delta}$, and $\mathcal{Q}(\widehat{\lambda}_t) = \mathcal{Q}(\lambda_{t-1})$. Hence $\mathcal{Q}(\widehat{\lambda}_t^k) = \mathcal{Q}(\lambda_{t-1}^k)$ on both paths. *The display.* By Proposition 1, $\mathrm{APD}_t^k = [X\,\mathcal{Q}(\widehat{\lambda}_t^k)^{\alpha}]^{1/\Gamma}$ with $X = (A_S^a/A_N^a)(A_S^n/A_N^n)^{\alpha+\gamma\varepsilon}$ the same on both paths, and the inherited ratios are $\lambda_{t-1}^D = \lambda_{-1}/(1+\xi)^t$ and $\lambda_{t-1}^U = T^t(\lambda_{-1})$. *Monotonicity and the limit.* The function $\mathcal{Q}$ is nondecreasing, λ_{t-1}^U is nondecreasing in t with limit χ, and λ_{t-1}^D is decreasing in t, so the numerator is nondecreasing, the denominator nonincreasing, and $\mathcal{A}_t$ nondecreasing; as $t \to \infty$ the denominator reaches $e^{-\delta}$, since $\lambda_{t-1}^D \to 0$, and the numerator converges to $\mathcal{Q}(\chi) = \chi$, since $\chi > e^{-\delta}$. *The case $\lambda_{-1} < e^{-\delta}$.* Both inherited ratios lie below $e^{-\delta}$ for $t < T^*(\lambda_{-1})$, so both factors equal $e^{-\delta}$ and $\mathcal{A}_t = 1$; from $t = T^*(\lambda_{-1})$ on, $\lambda_{t-1}^U \geq e^{-\delta}$ while $\lambda_{t-1}^D < e^{-\delta}$, which gives the stated value; it equals one only if $T^t(\lambda_{-1}) = e^{-\delta}$, which can happen at $t = T^*(\lambda_{-1})$ alone, since $T^t(\lambda_{-1})$ is strictly increasing in t toward $\chi > e^{-\delta}$. *Coincident paths.* If type-S innovation occurs in every period under laissez-faire, the two paths are the same sequence of decisions from the same start, so every object coincides. □

A.8 The Constrained Planner's Problem

This section defines the equilibrium of a period at an imposed innovation decision, the object that the constraint of Definition 2 refers to, and derives the recursive formulation of Section 3.1: the price index and consumption, the factorization (A.15), the treatment of the trend for every σ, the existence and uniqueness of the value function, the one-step policy-improvement argument behind the hypothesis of Proposition 4(iii), and the decomposition of the welfare difference between two paths that the proofs of Section A.9 use.

Definition A.1. Fix a period t, the inherited quality ratio $\lambda_{t-1} \in (0, \chi]$, and a decision $m_{S,t} \in \{0,1\}$. The *equilibrium of period t at the decision $m_{S,t}$* consists of the vintages, the prices of intermediate inputs and of agricultural goods, wages, farmers' choices of labor, input type, vintage, and input quantities, consumption allocations, and incomes of period t that satisfy the within-period conditions of Definition 1, with $m_{N,t} = 1$ and with the decision $m_{S,t}$ in place of the entry condition (12) for type-S inputs; the vintages are $q_{N,t} = (1+\xi)\,q_{N,t-1}$ and $q_{S,t}$ given by (10) at the decision $m_{S,t}$.

Lemma A.2 shows that this equilibrium exists and is unique for every $(\lambda_{t-1}, m_{S,t})$, and that the equilibrium path of Definition 1 is the sequence of these equilibria at the decisions $m_{S,t} =$

$\mathbb{1}\{\Pi_{S,t}(\lambda_{t-1}) \geq 0\}$. The planner's problem of Definition 2 replaces the entry condition for type-S inputs by the planner's decision and leaves every other condition of the period in place.

The price index and consumption. With the preferences (14) and $\mu = 1/2$, the household in country j spends the share ϕ of its income on agricultural goods and the share $1-\phi$ on the non-agricultural good, and it divides its expenditure on agricultural goods between the two varieties so that $c^a_{j,N,t}/c^a_{j,S,t} = (p^a_{N,t}/p^a_{S,t})^{-\varepsilon}$. The minimum expenditure that yields one unit of the agricultural aggregate $C^a_{j,t}$ is the price index of agricultural goods

$$P^a_t = \left[\tfrac{1}{2}\left(p^a_{N,t}\right)^{1-\varepsilon} + \tfrac{1}{2}\left(p^a_{S,t}\right)^{1-\varepsilon}\right]^{\frac{1}{1-\varepsilon}}, \tag{A.14}$$

which is the same for households in both countries, so $C^a_{j,t} = \phi E_{j,t}/P^a_t$ and $C^n_{j,t} = (1-\phi)E_{j,t}$. Substituting into (14),

$$C_{j,t} = \left(\frac{\phi E_{j,t}}{P^a_t}\right)^{\phi}\left((1-\phi)E_{j,t}\right)^{1-\phi} = \frac{E_{j,t}}{P_t}, \qquad P_t = \phi^{-\phi}(1-\phi)^{-(1-\phi)}\left(P^a_t\right)^{\phi},$$

and P_t is the minimum expenditure per unit of the consumption aggregate, since the aggregate is linearly homogeneous. Computing incomes and the price index of agricultural goods is therefore computing consumption.

The factorization of consumption. Fix a period t, the inherited quality ratio λ_{t-1}, and the decision $m_{S,t}$. By Proposition A.1(ii)–(vi), the input choices, the share $s_{S,t}$, the relative price of agricultural goods, world agricultural revenue, the revenues of the two countries, profits, and incomes depend on $(\lambda_{t-1}, m_{S,t})$ alone and not on $q_{N,t-1}$: the price-adjusted quality ratio $\widehat{\lambda}_t$ depends on the quality ratio only, the margin rates on $G_{N,t}$ and $G_{S,t}$, and the fixed costs are c and $2c$. By Proposition A.1(iii), the price-adjusted qualities of the inputs in use are $h^*_{N,t} = q_{N,t-1}/\psi$ and $h^*_{S,t} = \mathcal{Q}(\widehat{\lambda}_t)\,h^*_{N,t}$, both proportional to $q_{N,t-1}$. Solving (A.1) for the agricultural output price,

$$p^a_{j,t} = \frac{R^{\gamma}_{j,t}}{\alpha^{\alpha}\beta^{\beta}A^a_j(A^n_j)^{-\beta}\left(h^*_{j,t}\right)^{\alpha}} = \tilde{p}^a_j(\lambda_{t-1}, m_{S,t})\left(q_{N,t-1}\right)^{-\alpha},$$

with $\tilde{p}^a_j$ the price at a unit level of frontier quality. The index P^a_t is linearly homogeneous in the two prices, so $P^a_t = \tilde{P}^a(\lambda_{t-1}, m_{S,t})\left(q_{N,t-1}\right)^{-\alpha}$ and $P_t = \tilde{P}(\lambda_{t-1}, m_{S,t})\left(q_{N,t-1}\right)^{-\alpha\phi}$. Since the frontier technology improves in every period and $q_{N,-1} = 1$, $q_{N,t-1} = (1+\xi)^t$, and

$$C_{j,t} = \frac{E_{j,t}}{P_t} = (1+\xi)^{\alpha\phi t}\,\widetilde{C}_j(\lambda_{t-1}, m_{S,t}), \qquad \widetilde{C}_j := \frac{E_j}{\tilde{P}}, \qquad j \in \{N,S\}, \tag{A.15}$$

which is the factorization that Section 3.1 states: $\widetilde{C}_j(\lambda, m_S)$ is the consumption aggregate of the household in country j in the equilibrium of a period with inherited quality ratio λ and decision m_S, evaluated at a unit level of frontier quality. The factorization is a statement about one period, so it holds along every path.

The trend for every σ. Let $g \coloneqq (1+\xi)^{\alpha\phi}$. For $\sigma \neq 1$, $u(C_{j,t}) = [(g^t \widetilde{C}_{j,t})^{1-\sigma} - 1]/(1-\sigma) = g^{(1-\sigma)t}\, u(\widetilde{C}_{j,t}) + (g^{(1-\sigma)t} - 1)/(1-\sigma)$, and for $\sigma = 1$, $\ln C_{j,t} = \ln \widetilde{C}_{j,t} + t \ln g$. In both cases the second term does not depend on the sequence of decisions, and its discounted sum is finite by Assumption A.4(d). The criterion (24) is therefore, up to a constant common to every sequence of decisions,

$$\sum_{t=0}^{\infty} \beta_\sigma^t \mathcal{W}(\lambda_{t-1}, m_{S,t}), \qquad \beta_\sigma \coloneqq \frac{(1+\xi)^{\alpha\phi(1-\sigma)}}{1+\rho} < 1, \tag{A.16}$$

with $\mathcal{W}(\lambda, m_S) = v_N\, u(\widetilde{C}_N(\lambda, m_S)) + v_S\, u(\widetilde{C}_S(\lambda, m_S))$, which is the problem behind (25); at $\sigma = 1$, $\beta_1 = 1/(1+\rho)$. Because the trend is removed exactly, the optimal sequence of decisions of (A.16) is the optimal sequence of (24). Propositions 2, 3, and 5 and Lemmas A.5 and A.6 are stated and proved for general σ with β_σ in place of $1/(1+\rho)$; Section 3 displays the case $\sigma = 1$.

Lemma A.5. *The per-period welfare function $\mathcal{W}$ is bounded on $(0, \chi] \times \{0,1\}$. The operator $(\mathcal{T}V)(\lambda) = \max_{m_S \in \{0,1\}} \{\mathcal{W}(\lambda, m_S) + \beta_\sigma V(\lambda'(\lambda, m_S))\}$ is a contraction of modulus β_σ on the space of bounded functions on $(0, \chi]$ with the supremum norm. The value function V is its unique fixed point and the limit of value iteration from any bounded function; it is the value of the problem (A.16) over all sequences of decisions, and the rule $m_S(\lambda) = \mathbb{1}\{\Pi^{SP}(\lambda) \geq 0\}$, with Π^{SP} defined by (26) with β_σ, attains it from every λ_{-1}. The function V can be discontinuous at $e^{-\delta}$, where $\mathcal{W}$ jumps, and at the preimages of $e^{-\delta}$ under the two transition maps.*

Proof. Both transition maps send $(0, \chi]$ into itself: T is increasing with $T((0,\chi]) = (a\chi, \chi]$, and $\lambda/(1+\xi) \in (0, \chi]$. Incomes are strictly positive by Assumption A.4(a)–(c), and every object of the period is a continuous function of λ on each of $(0, e^{-\delta})$ and $[e^{-\delta}, \chi]$ that is constant on the first interval (Proposition A.1: below $e^{-\delta}$ the input choices and every ratio are the same for every λ) and continuous on the compact second interval, so $\mathcal{W}$ is bounded. On the space of bounded functions with the supremum norm, $\mathcal{T}$ maps bounded functions to bounded functions, is monotone, and satisfies $\mathcal{T}(V + k) = \mathcal{T}V + \beta_\sigma k$ for every constant k, so it is a contraction by Blackwell's sufficient conditions, with a unique fixed point approached by value iteration from any bounded function. With bounded per-period welfare and $\beta_\sigma < 1$, the value of the sequence problem (A.16) is the unique bounded fixed point of $\mathcal{T}$, and a rule that attains the maximum in (25) at every λ attains the value from every start; the rule $\mathbb{1}\{\Pi^{SP} \geq 0\}$ does so by the definition of Π^{SP}, with ties resolved toward innovation. The possible discontinuities come from the jump of $\mathcal{W}$ at $e^{-\delta}$, where Southern farmers switch input types, and from the continuation terms, which can jump where the next period's quality ratio crosses $e^{-\delta}$; nothing in the argument uses continuity. □

Lemma A.6. *Let $V^1(\lambda)$ denote the value of (A.16) from λ under the rule that improves type-S inputs in every period, and let*

$$J(\lambda) \coloneqq \Big[\mathcal{W}(\lambda, 1) + \beta_\sigma V^1\big(T(\lambda)\big)\Big] - \Big[\mathcal{W}(\lambda, 0) + \beta_\sigma V^1\Big(\frac{\lambda}{1+\xi}\Big)\Big]$$

be the gain from improving type-S inputs now rather than one period later when they are improved in every later period either way.

(a) If $J(\lambda) \geq 0$ for every $\lambda \in (0,\chi]$, the rule that improves type-S inputs in every period is optimal, $V = V^1$, and $\Pi^{SP} = J$.

(b) For a fixed sequence of decisions, the criterion is linear in the weights, so $J = v_N J_N + v_S J_S$, where J_j is the gain of part (a) computed with country j's lifetime utility alone. If $J_N(\lambda) > 0$ and $J_S(\lambda) > 0$ for every λ, the rule is optimal for every weighting (v_N, v_S), and $\Pi^{SP}(\lambda) = J(\lambda) > 0$ for every weighting.

Proof. (a) V^1 is the unique bounded fixed point of the operator $V \mapsto \mathcal{W}(\cdot,1) + \beta_\sigma V(T(\cdot))$, a contraction by the argument of Lemma A.5. If $J \geq 0$ everywhere, the decision $m_S = 1$ attains $\max_{m_S}\{\mathcal{W}(\lambda,m_S) + \beta_\sigma V^1(\lambda'(\lambda,m_S))\}$ at every λ, so $\mathcal{T}V^1 = V^1$; by the uniqueness of the fixed point of $\mathcal{T}$, $V^1 = V$, the rule is optimal, and Π^{SP}, defined with $V = V^1$, equals J. (b) The per-period welfare function is $\mathcal{W} = v_N u(\widetilde{C}_N) + v_S u(\widetilde{C}_S)$, so for a fixed sequence of decisions the criterion, and hence V^1 and J, are the same weighted sums of the countries' objects. If $J_N > 0$ and $J_S > 0$ everywhere, then $J > 0$ everywhere for every weighting, and part (a) applies to each weighting. □

Lemma A.7. *Fix $\lambda \in (0,\chi]$ and a decision rule g for the periods after the current one, a function from $(0,\chi]$ to $\{0,1\}$, and let $V^g(\lambda)$ denote the value of (A.16) from λ when every decision follows g. Compare the path on which type-S inputs are improved in the current period with the path on which they are not, both following g afterwards. The difference in welfare (24) between the two paths equals*

$$\big[\mathcal{W}(\lambda,1) - \mathcal{W}(\lambda,0)\big] + \beta_\sigma\Big[V^g\big(T(\lambda)\big) - V^g\Big(\frac{\lambda}{1+\xi}\Big)\Big],$$

the current-period effect plus the continuation value under g. With g the planner's optimal rule, $V^g = V$ and the difference is the social value $\Pi^{SP}(\lambda)$ of (26), with β_σ in place of $1/(1+\rho)$.

Proof. V^g is the unique bounded fixed point of $V \mapsto \mathcal{W}(\cdot,g(\cdot)) + \beta_\sigma V(\lambda'(\cdot,g(\cdot)))$, a contraction by the argument of Lemma A.5. On the first path the decision in the current period is 1 and the next period's inherited ratio is $T(\lambda)$; on the second the decision is 0 and the next period's ratio is $\lambda/(1+\xi)$; from the next period on both paths follow g from their own ratios. The difference in (A.16) between the two paths is the difference in the current period, $\mathcal{W}(\lambda,1) - \mathcal{W}(\lambda,0)$, plus $\beta_\sigma[V^g(T(\lambda)) - V^g(\lambda/(1+\xi))]$, since every later term depends only on the path's own ratio and decision. The trend cancels term by term, so the same holds for (24). With g the planner's optimal rule, $V^g = V$ and the right-hand side is (26). □

A.9 Proofs of Propositions 2–4 and Corollary 1

Proof of Proposition 2. The firm's payoff is its operating profit in the period less the fixed cost, $\pi_{S,t}(\lambda_{t-1}) - c$, plus any subsidy it receives in the period; the operating profit is a function of

the period's objects (Proposition A.1(vi)), which the equilibrium of the period at the decision determines (Definition A.1 and Lemma A.2). The payoff contains no object of a later period, for three reasons. The patent expires at the end of the period, so the vintage yields the firm no revenue afterwards. The blueprint is disclosed at expiry and the vintage is then supplied at marginal cost (Assumption A.2(a)), so the vintage has no resale value. And the firm exits at the end of the period with no claim on later research opportunities and no counterparty with which to contract (Assumption A.3(d)). The payoff is therefore invariant to every term of the continuation value of Lemma A.7, whatever the decision rule of the later periods. □

Proof of Proposition 3. The private value. By Corollary A.1, for $\lambda_{t-1} < e^{-\delta}$, $\widehat{\lambda}_t < e^{-\delta}$ whether or not type-S innovation occurs (Proposition A.1(ii) under Assumption A.5(a)), Southern farmers use type-N inputs, $\pi_{S,t}(\lambda_{t-1}) = 0$, and $\Pi_{S,t}(\lambda_{t-1}) = -c$, so the entry condition (12) fails.

The two configurations of the period. In both configurations $\widehat{\lambda}_t < e^{-\delta}$, so by Proposition A.1(ii)–(iv) both have $\omega_{S,S,t} = 0$, $\mathcal{Q}(\widehat{\lambda}_t) = e^{-\delta}$, the same $\mathcal{F}_t$, the same relative price, the same share $s_{S,t} = s_S^0 := [1 + \mathcal{F}_t^{-(\varepsilon-1)/\Gamma}]^{-1}$, and, by Proposition 1, the same agricultural productivity differences; none of these objects depends on λ_{t-1}. Type-S inputs have no market, so $\kappa_t = \kappa_0 := 1 - \phi\gamma - \phi\alpha\nu_N$ in both configurations, and by (A.5), $R_t^W = \phi B/\kappa_0$ with $B := A_N^n L_N + A_S^n L_S - c\,(1 + m_{S,t})$: the two configurations differ only through B, that is, through the fixed cost.

The social value. By Lemma A.7 at the planner's rule, $\Pi^{\mathrm{SP}}(\lambda_{t-1})$ is the current-period effect plus the continuation value. The current-period effect is the change in $\nu_N \ln \widetilde{C}_N + \nu_S \ln \widetilde{C}_S$ when B falls from $A_N^n L_N + A_S^n L_S - c$ to $A_N^n L_N + A_S^n L_S - 2c$ at the fixed ratios of the previous step; nothing in it depends on λ_{t-1}. The social value is therefore the continuation value net of the welfare cost of the fixed cost, which is the same at every $\lambda_{t-1} < e^{-\delta}$, and the planner improves type-S inputs if and only if the continuation value is at least that cost. Neither the private value nor the social value contains anything that the firm could appropriate. □

Remark: the sign of the current-period effect for each household. Let $K := \phi/\kappa_0$ and $W_S := A_S^n L_S$. Below $e^{-\delta}$ the frontier patent holder serves every farmer, so $\pi_{N,t} = \nu_N \alpha R_t^W$, and by (16),

$$E_{N,t}(B) = -W_S + \big[1 + K\big(\gamma(1 - s_S^0) + \alpha\nu_N\big)\big]B, \qquad E_{S,t}(B) = W_S + K\gamma s_S^0 B,$$

using $A_N^n L_N = B + c(1 + m_{S,t}) - W_S$ and the net profits $\pi_{N,t} - c$ and, with innovation, $\pi_{S,t} - c = -c$. Both agricultural output prices are proportional to $R_{j,t}^{\gamma}$ at the fixed ratios, and $R_{j,t} \propto B$, so $P_t^a \propto B^{\gamma}$ and $\ln \widetilde{C}_{j,t} = \ln E_{j,t}(B) - \phi\gamma \ln B$ up to a constant. For the North, $d \ln \widetilde{C}_N / d \ln B = BE_N'(B)/E_N(B) - \phi\gamma > 1 - \phi\gamma > 0$ at every B, because $E_N(B) < BE_N'(B)$; so the fixed cost, which lowers B, lowers Northern consumption. For the South, $d \ln \widetilde{C}_S / d \ln B = \gamma R_{S,t}/E_{S,t} - \phi\gamma = \gamma\,(R_{S,t}/E_{S,t} - \phi)$, the share of farmers' profits in Southern income less $\phi\gamma$, and $R_{S,t}/E_{S,t}$ is increasing in B. If $R_{S,t} > \phi E_{S,t}$ in the configuration with innovation, the smaller value of B, then the derivative is positive on the whole interval and the fixed cost lowers Southern consumption too: the fixed cost lowers Southern farmers' profits by more than it lowers the price index of agricultural goods. If $R_{S,t} < \phi E_{S,t}$ in the configuration without innovation, the larger value of

B, the derivative is negative on the whole interval and Southern consumption rises. The fixed cost is paid by Northern households and reaches the South only through world agricultural revenue.

Lemma A.8. *Fix $\lambda_{t-1} < e^{-\delta}$, and suppose that type-S inputs are improved in every later period on both paths of Lemma A.7. Then consumption in both countries is the same on the two paths in every period s with $1 \le s < T^*(\lambda_{t-1})$ after the current one. In period $T^*(\lambda_{t-1})$, Southern farmers switch to type-S inputs on the path with the earlier innovation and not on the other, and from then on the lead of the earlier innovation in the quality ratio declines geometrically:*

$$T^s(\lambda_{t-1}) - T^{s-1}\Big(\frac{\lambda_{t-1}}{1+\xi}\Big) = (1-a)^{s-1}\Big[a\chi - \lambda_{t-1}\Big(\frac{1}{1+\xi} - (1-a)\Big)\Big], \qquad s \ge 1, \tag{A.17}$$

where T^s denotes the s-fold iterate of the catch-up map in (17); at $s = 1$ the lead is the catch-up $a(\chi - \lambda_{t-1})$ that the innovation achieves plus the decay $\lambda_{t-1}\xi/(1+\xi)$ that the other path suffers.

Proof. On the path with the innovation in the current period the inherited ratio in period $s \ge 1$ after the current one is $T^s(\lambda_{t-1})$; on the other path it is $T^{s-1}(\lambda_{t-1}/(1+\xi))$, which is smaller because T is increasing and $\lambda_{t-1}/(1+\xi) < \lambda_{t-1} \le T(\lambda_{t-1})$. For $1 \le s < T^*(\lambda_{t-1})$, both ratios lie below $e^{-\delta}$ (Section A.7), both configurations improve type-S inputs and pay $2c$, and by the second step of the proof of Proposition 3 the two equilibria coincide in every object, so consumption is the same on both paths. In period $s = T^*(\lambda_{t-1})$, the first path's ratio is at or above $e^{-\delta}$, so Southern farmers switch to type-S inputs there, while the second path's ratio is $T^{T^*-1}(\lambda_{t-1}/(1+\xi)) < T^{T^*-1}(\lambda_{t-1}) < e^{-\delta}$. Finally, T is affine with slope $1-a$, so $T^s(\lambda) - T^{s-1}(\lambda/(1+\xi)) = (1-a)^{s-1}[T(\lambda) - \lambda/(1+\xi)]$, and $T(\lambda) - \lambda/(1+\xi) = a\chi - \lambda[1/(1+\xi) - (1-a)] = a(\chi-\lambda) + \lambda\xi/(1+\xi)$, which is (A.17). □

Proof of Proposition 4. (i) is Corollary A.1 and the entry condition (12).

(ii) By (i), from $\lambda_{-1} < e^{-\delta}$ the laissez-faire path never improves type-S inputs, and its inherited ratio in period t is $\lambda_{t-1} = \lambda_{-1}/(1+\xi)^t$. Write β for β_σ and $V^{\mathrm{LF}}(\lambda)$ for the welfare (A.16) of the laissez-faire path from λ. Suppose first that $\Pi^{\mathrm{SP}}(\lambda_{t_0-1}) > 0$ for some $t_0 \ge 0$. The welfare of the laissez-faire path from λ_{-1} is $\sum_{t<t_0}\beta^t\mathcal{W}(\lambda_{t-1},0) + \beta^{t_0}V^{\mathrm{LF}}(\lambda_{t_0-1})$, and $V^{\mathrm{LF}}(\lambda_{t_0-1}) \le \mathcal{W}(\lambda_{t_0-1},0) + \beta V(\lambda_{t_0-1}/(1+\xi))$, since V is the maximal value from the next period's ratio; by $\Pi^{\mathrm{SP}}(\lambda_{t_0-1}) > 0$, this is smaller than $\mathcal{W}(\lambda_{t_0-1},1) + \beta V(T(\lambda_{t_0-1})) = V(\lambda_{t_0-1})$. The sequence that follows laissez-faire up to period $t_0 - 1$, improves type-S inputs in period t_0, and continues optimally is feasible for the planner and attains $\sum_{t<t_0}\beta^t\mathcal{W}(\lambda_{t-1},0) + \beta^{t_0}V(\lambda_{t_0-1})$, which is strictly larger. The laissez-faire sequence is therefore not optimal, so it is not a sequence that the planner chooses, and the planner's sequence yields strictly higher welfare. Suppose instead that $\Pi^{\mathrm{SP}}(\lambda_{t-1}) < 0$ for every $t \ge 0$. Then at every ratio of the laissez-faire path the decision 0 is the unique maximizer in (25), so $V(\lambda_{t-1}) = \mathcal{W}(\lambda_{t-1},0) + \beta V(\lambda_t)$ for every t, and by induction $V(\lambda_{-1}) = \sum_{t<n}\beta^t\mathcal{W}(\lambda_{t-1},0) + \beta^n V(\lambda_{n-1})$ for every n; since V is bounded and $\beta < 1$, letting $n \to \infty$ gives $V(\lambda_{-1}) = V^{\mathrm{LF}}(\lambda_{-1})$. The laissez-faire sequence attains the planner's

value, and since the maximizer is unique at every ratio it visits, it is the planner's sequence of decisions.

(iii) By Lemma A.5, the rule $\mathbb{1}\{\Pi^{\mathrm{SP}}(\lambda) \geq 0\}$ is optimal, and under the hypothesis it equals 1 at every λ; the planner's path is $\lambda_t = T^{t+1}(\lambda_{-1})$, which rises to χ by the closed form of Proposition A.2. The laissez-faire sequence of decisions is feasible for the planner. Let t_0 be the first period in which laissez-faire does not improve type-S inputs; it exists whenever the laissez-faire path reaches a ratio at which $\Pi_{S,t} < 0$, and $t_0 = 0$ from $\lambda_{-1} < e^{-\delta}$ by (i). Before t_0 the two paths coincide, and at t_0 the argument of (ii), with $\Pi^{\mathrm{SP}}(\lambda_{t_0-1}) > 0$ by the hypothesis, shows that the laissez-faire path yields strictly lower welfare than the planner's from λ_{-1}: the equilibrium is constrained inefficient. If type-S innovation occurs in every period under laissez-faire, the laissez-faire sequence is the planner's optimal sequence, so the equilibrium is constrained efficient. □

Proof of Corollary 1. The hypothesis for the polar planner of country $j \in \{N, S\}$ is $\Pi_j^{\mathrm{SP}}(\lambda) > 0$ at every λ, where Π_j^{SP} is the social value computed with country j's lifetime utility alone. By Lemma A.5, the rule $\mathbb{1}\{\Pi_j^{\mathrm{SP}} \geq 0\}$ is optimal for that planner, so the rule that improves type-S inputs in every period is optimal, its value is $V_j = V_j^1$, and $\Pi_j^{\mathrm{SP}} = J_j > 0$. By Lemma A.6(b), the rule is optimal for every weighting. From $\lambda_{-1} < e^{-\delta}$, the argument of the proof of Proposition 4(ii) applied to each country's criterion, with $t_0 = 0$ and $\Pi_j^{\mathrm{SP}}(\lambda_{-1}) > 0$, shows that each household's lifetime utility is strictly higher on the planner's path than on the laissez-faire path. □

A.10 The R&D Subsidy and Longer Patents

This section introduces the R&D subsidy, shows that it implements the planner's decisions, and proves Proposition 5 in the extended game with longer patents, where it also shows that on the laissez-faire path the length of the patent changes nothing above $e^{-\delta}$ where laissez-faire sustains type-S innovation (Proposition A.5). The baseline model of Section 2 has no government; the quantitative model of Section 6 introduces the subsidy at the level of product lines, and the subsidy below is its counterpart with one product line per input type.

The subsidy. The government pays a subsidy $s_t \geq 0$ to the R&D firm that improves type-S inputs in period t, so that the firm's entry condition becomes $\pi_{S,t} + s_t \geq c$, and finances it with lump-sum taxes $\tau_{N,t}$ and $\tau_{S,t}$ on the two households, with $\tau_{N,t} + \tau_{S,t} = s_t m_{S,t}$. Household incomes are

$$E_{N,t} = E_{N,t}^{\mathrm{pre}} + s_t m_{S,t} - \tau_{N,t}, \qquad E_{S,t} = E_{S,t}^{\mathrm{pre}} - \tau_{S,t}, \tag{A.18}$$

where $E_{j,t}^{\mathrm{pre}}$ are the incomes (16) and the subsidy reaches Northern households as the profit of their R&D firm. I consider three financing conventions: Northern financing, $\tau_{N,t} = s_t m_{S,t}$; Southern financing, $\tau_{S,t} = s_t m_{S,t}$; and proportional financing, $\tau_{j,t} = s_t m_{S,t} E_{j,t}^{\mathrm{pre}} / (E_{N,t}^{\mathrm{pre}} + E_{S,t}^{\mathrm{pre}})$. The subsidy and the taxes cancel in world income, so world agricultural revenue, the revenues

of the two countries, prices, and every object of Proposition A.1 are unchanged by the subsidy given the innovation decisions: the subsidy changes the entry decision and the division of income across the two households, and nothing else within the period. Definition 1 extends by the requirements that entry satisfies the subsidized condition, that households optimize at the incomes (A.18), and that the government's budget balances; under Southern and proportional financing I require incomes to remain positive.

Proposition A.4. *Suppose that the hypothesis of Proposition 4(iii) holds, and let*

$$\bar{s}(\lambda) := \max\{0,\ -\Pi_{S,t}(\lambda)\} = \max\{0,\ c - \pi_{S,t}(\lambda)\}, \tag{A.19}$$

which equals c for every $\lambda < e^{-\delta}$ and zero wherever type-S innovation is profitable under laissez-faire.

(i) Under each financing convention, the subsidy $s_t = \bar{s}(\lambda_{t-1})$ makes the subsidized entry condition hold at every quality ratio, so type-S inputs are improved in every period from every λ_{-1}: the decentralized equilibrium with the subsidy follows the planner's path.

(ii) Under Northern financing, the subsidy and the tax cancel within the Northern household, incomes equal the incomes (16) of the planner's path in every period, and so does consumption.

(iii) Under Southern and proportional financing, the path of the quality ratio, prices, and revenues is the same, and only the division of consumption across the two households differs from the planner's path.

(iv) $\bar{s}$ is the smallest subsidy that implements the planner's path: any smaller payment fails the entry condition at some quality ratio below $e^{-\delta}$, and any larger payment changes no price, revenue, or innovation decision and is a transfer between the households.

(v) If type-S innovation is profitable at every $\lambda \in [e^{-\delta}, \chi]$ (case (a) of Proposition A.2), then from $\lambda_{-1} < e^{-\delta}$ the subsidy is paid in exactly $T^*(\lambda_{-1})$ consecutive periods, each at the full cost c, and never afterwards; its discounted cost is $c\sum_{t=0}^{T^*(\lambda_{-1})-1}(1+\rho)^{-t}$.

Proof. (i) At $\lambda < e^{-\delta}$, $\pi_{S,t}(\lambda) + \bar{s}(\lambda) = 0 + c = c$, so the subsidized condition holds and an indifferent firm enters; at $\lambda \geq e^{-\delta}$ it holds with $\bar{s}(\lambda) = \max\{0, c - \pi_{S,t}(\lambda)\}$ by construction. The subsidy is paid to the firm whichever household is taxed, so the entry decisions, and with them the path, are the same under every convention, and they are the planner's decisions under the hypothesis of Proposition 4(iii). (ii) Under Northern financing, $E_{N,t} = E_{N,t}^{\text{pre}}$ and $E_{S,t} = E_{S,t}^{\text{pre}}$ in every period, and the planner's path has the same decisions and no fiscal flows, so incomes coincide and, by the factorization of Appendix A.8, so does consumption. (iii) The subsidy and the taxes cancel in world income, so world agricultural revenue, the revenues, and prices are unchanged given the decisions; the conventions differ only in which household pays. (iv) Below $e^{-\delta}$ any payment smaller than c leaves $\pi_{S,t} + s < c$ and no entry; any payment larger than $\bar{s}$ leaves the decision, the quality ratio, and every price and revenue unchanged, and moves only

the division of income between the households. (v) From $\lambda_{-1} < e^{-\delta}$ every period's inherited ratio is $T^t(\lambda_{-1})$, which lies below $e^{-\delta}$ exactly for $t < T^*(\lambda_{-1})$ (the definition of T^* in Section A.7); the subsidy is c at those dates and zero afterwards, because the ratio then lies in $[e^{-\delta}, \chi]$, which T maps into itself and on which $\Pi_{S,t} \geq 0$ in case (a). □

Three remarks are in order. First, the subsidy is a transfer that flips the sign of the firm's decision, not a price of the externality: below $e^{-\delta}$ it equals the fixed cost, whatever the continuation value that the innovation creates, and any larger payment implements the same path. This differs from the Pigouvian tax of the pecuniary-externality literature, which is set equal to the uninternalized effect (Bianchi, 2011; Jeanne & Korinek, 2019); here the decision is binary, and policy is judged by the decisions it changes rather than by the size of the wedge. Second, a subsidy paid for fewer than $T^*(\lambda_{-1})$ periods from $\lambda_{-1} < e^{-\delta}$ and then withdrawn changes no allocation at any date: the quality ratio never reaches $e^{-\delta}$, Southern farmers never switch, every object of every period is the one without innovation, and the only effect is the fixed cost paid in the program periods, which lowers Northern consumption in each of them, by the remark after the proof of Proposition 3. The gain from a program of type-S innovations arrives with the innovation that lifts the quality ratio to $e^{-\delta}$, and a program of $T^*(\lambda_{-1})$ periods is enough in case (a), because the decentralized equilibrium sustains type-S innovation on its own from then on. Third, Section 6 evaluates a subsidy that Southern households finance. The corresponding planner is the Southern planner of Definition 2 with Southern income net of the tax, $E_{S,t} = E^{\text{pre}}_{S,t} - \bar{s}(\lambda_{t-1}) m_{S,t}$, in every period in which it improves type-S inputs. By part (iii), the path of the quality ratio, prices, and revenues at a given sequence of decisions is unchanged, so this planner's problem differs from Definition 2 only through the per-period welfare function, which in every period with type-S innovation includes the welfare cost to Southern households of the subsidy they pay; its social value of type-S innovation is the South's continuation value, itself net of the subsidies of later periods, net of the cost of the current subsidy, and its hypothesis is more demanding than the Southern planner's of Section 3.2, where Northern households bear the cost.

Longer patents: the extended game, the proof of Proposition 5, and patent length above $e^{-\delta}$. Let patents last $T_p \geq 1$ periods. The R&D firm that improves type-k inputs in period t receives the exclusive right to produce and sell its vintage in periods $t, \ldots, t + T_p - 1$; at the end of period $t + T_p - 1$ the patent expires and the blueprint becomes public, as in Assumption A.2(a) applied at T_p. The research opportunity of each period still belongs to a new firm, which cannot contract with the holders of earlier patents (Assumption A.3(d)), so the firm that improves type-k inputs in period t competes in prices, within Stage 2 of each later period of its patent, with the holders of the newer vintages of the same input type and with the competitive fringe, which supplies the newest disclosed vintage at marginal cost ψ. The timing of Assumption A.2(b) applies in every period: farmers commit to an input type before the prices of the period are posted, so a patent holder cannot attract farmers who chose the other input type by lowering its price,

and it sets its price against the alternatives of its own input type. As in Lemma A.1, each farmer buys the vintage of her input type with the highest quality per unit of expenditure, and a farmer who is indifferent between two vintages of her input type buys the newer one, as in Assumption A.2(b); a holder whose vintage loses to a newer vintage of the same type offers it at marginal cost, and the holder of a vintage of an input type that no farmer chose posts the price of Lemma A.1. Farmers anticipate, for each input type, the price that its holders would post if they chose that type; a lower price that a holder would not post once farmers have chosen its type, such as marginal cost, does not enter their choice. A firm enters in period t if and only if the discounted operating profits of its patent over its life, given the entry decisions of the firms of later periods, cover the fixed cost c; the case $T_p = 1$ is the model of Section 2. The patent status of the vintages inherited at date 0 plays no role in what follows: below $e^{-\delta}$ no farmer buys type-S inputs, and along a path on which type-S inputs are improved in every period, as under the hypothesis of Proposition A.5, a newer vintage is created in every period, so an inherited vintage is offered at ψ either as disclosed or as losing.

Proof of Proposition 5. Consider the type-S patent created in period t and a period $t+n$ with $0 \leq n < T_p$. *The profit of a patent by age.* If a newer type-S vintage exists in period $t+n$, the holder of the newest one sets its price against the best alternative of its type at marginal cost, and a farmer who is indifferent buys the newer vintage, so the older vintage attracts no demand at any price at or above marginal cost, its holder posts ψ, and it earns nothing. If no newer vintage exists but Southern farmers use type-N inputs in period $t+n$, no farmer buys type-S inputs, since Northern farmers never do (Proposition A.1(ii)). Hence the patent earns operating profit only in periods in which Southern farmers use type-S inputs and its vintage is the newest type-S vintage.

(i) Suppose that $\lambda_{t-1} < e^{-\delta}$ and that the patent has not been superseded by a newer vintage of type-S inputs in period $t+n$. The best alternative offer of type-S inputs is the previous vintage, of base quality $q_{S,t-1}$, at price ψ: supplied by the competitive fringe if its patent has expired, and offered at marginal cost by its holder if it has not, because that vintage attracts no demand against the newer one. The argument of Lemma A.1 therefore applies in every period of the patent's life, and the holder posts the price of (7), $p_S^x = \min\{\alpha, G_{S,t}\,\psi\}$, in each of them. Its price-adjusted quality $q_{S,t}/p_S^x$ is constant over the patent's life, while that of type-N inputs is $q_{N,t+n-1}/\psi$ in period $t+n$ by Proposition A.1(i), since the frontier technology improves in every period and its innovation is non-drastic; it grows by the factor $1+\xi$ per period. The ratio of the two in period $t+n$ is

$$\frac{q_{S,t}/p_S^x}{q_{N,t+n-1}/\psi} = \frac{q_{S,t}/p_S^x}{q_{N,t-1}/\psi} \cdot \frac{1}{(1+\xi)^n} = \frac{\widehat{\lambda}_t}{(1+\xi)^n},$$

where $\widehat{\lambda}_t = \max\{\lambda_{t-1}, \alpha(1+\xi)T(\lambda_{t-1})\}$ is the price-adjusted quality ratio of period t (Proposition A.1(i)). Under Assumption A.5(a), $\widehat{\lambda}_t < e^{-\delta}$ strictly: on $[\lambda^{\mathrm{D}}, e^{-\delta})$ the innovation is non-drastic and $\widehat{\lambda}_t = \lambda_{t-1} < e^{-\delta}$; on $(0, \lambda^{\mathrm{D}})$ it is drastic and $\widehat{\lambda}_t = \alpha(1+\xi)T(\lambda_{t-1})$, which is in-

creasing in λ_{t-1} if $a < 1$ and equals λ^{D} at $\lambda_{t-1} = \lambda^{\mathrm{D}}$, so $\widehat{\lambda}_t < \lambda^{\mathrm{D}} \leq e^{-\delta}$, and if $a = 1$ it equals $\lambda^{\mathrm{D}} < e^{-\delta}$. Hence the ratio lies strictly below $e^{-\delta}$ at every age, and by the comparison in the proof of Proposition A.1(ii), $h_{S,S,t+n} < h_{S,N,t+n}$: no Southern farmer adopts the patented vintage, and no Northern farmer ever does. The patent's operating profit is zero in every period of its life whether or not it is superseded, since a superseded vintage earns nothing either. Entry in period t therefore yields $-c$ for every T_p, and type-S innovation does not occur at any $\lambda_{t-1} < e^{-\delta}$ for any patent length.

(ii) *The first-period profit.* In period t the vintage is the newest. The best alternative offer of type-S inputs is the inherited vintage $q_{S,t-1}$ at ψ, supplied by the competitive fringe if its patent has expired and offered at marginal cost by its holder otherwise, since it loses to the newer vintage. Frontier innovation occurs in period t for every T_p: the frontier firm's first-period operating profit covers its fixed cost whether or not type-S innovation occurs (Assumption A.3(b)), and its later profits are nonnegative. The holder of the new type-S vintage therefore posts the price of Lemma A.1, the price-adjusted quality ratio of the period is $\widehat{\lambda}_t$ of Proposition A.1(i), and the equilibrium of the period is that of Definition A.1 at $m_{S,t} = 1$. The patent's first-period operating profit is $\pi_{S,t}(\lambda_{t-1})$ of (A.7), which is zero if $\lambda_{t-1} < e^{-\delta}$. *An innovation followed by another.* If type-S inputs are improved again in period $t+1$, a newer vintage exists in every period $t+n$ with $n \geq 1$, so the patent earns $\pi_{S,t}(\lambda_{t-1})$ in period t and nothing afterwards, and entry in period t yields $\pi_{S,t}(\lambda_{t-1}) - c$ for every T_p. *Innovation in every period.* Suppose that, under patents of duration T_p, type-S inputs are improved in every period along the laissez-faire path from λ_{-1}. Every innovation is then followed by another in the next period, so the firm that enters in period t earns $\pi_{S,t}(\lambda_{t-1}) - c$, which is nonnegative because the firm enters. The entry condition (12) therefore holds at every quality ratio $\lambda_{t-1} = T^t(\lambda_{-1})$ of this path, and type-S inputs are improved in every period along the laissez-faire path from λ_{-1} under one-period patents as well. □

The length of the patent also changes nothing above $e^{-\delta}$ where laissez-faire sustains type-S innovation.

Proposition A.5. *In the game of this section, suppose that $\pi_{S,t}(\lambda) \geq c$ for every $\lambda \in [e^{-\delta}, \chi]$, so that under one-period patents type-S innovation occurs in every period in which Southern farmers use type-S inputs. Then, for every T_p, type-S inputs are improved in every period in which $\lambda_{t-1} \geq e^{-\delta}$, whatever the entry decisions of later firms, and every type-S patent is superseded one period after its creation. From every initial quality ratio, the laissez-faire path of innovation decisions, prices, allocations, and incomes is the same as under one-period patents.*

Proof. The first-period profit. By the proof of Proposition 5(ii), the first-period operating profit of a type-S patent created at $\lambda_{t-1} \geq e^{-\delta}$ is $\pi_{S,t}(\lambda_{t-1})$, which covers c by hypothesis. *Dominant entry.* Entry within a period is sequential and observable (Stage 1 of Section 2.1.2), so the two entry decisions of a period are determined by backward induction. The frontier firm's first-period operating profit covers its fixed cost whether or not type-S innovation occurs in the period

(Assumption A.3(b)); given frontier innovation, the type-S firm's first-period operating profit at $\lambda_{t-1} \geq e^{-\delta}$ is $\pi_{S,t}(\lambda_{t-1}) \geq c$; the operating profits of either patent in the later periods of its life are nonnegative; and not entering yields zero. Both firms therefore enter, whatever the entry decisions of the firms of later periods, an indifferent firm entering by the tie convention of (12). Hence type-N innovation occurs in every period for every T_p, type-S innovation occurs at every $\lambda_{t-1} \geq e^{-\delta}$ and, by Proposition 5(i), at no $\lambda_{t-1} < e^{-\delta}$; these are the decisions of the one-period model at every state, since under the hypothesis $\Pi_{S,t}(\lambda) \geq 0$ exactly on $[e^{-\delta}, \chi]$ (Corollary A.1). *The path.* Since T maps $[e^{-\delta}, \chi]$ into itself, from $\lambda_{-1} \geq e^{-\delta}$ a new type-S vintage is created in every period, so each patent is superseded one period after its creation; its holder then posts ψ and sells nothing (the proof of Proposition 5). From $\lambda_{-1} < e^{-\delta}$, no type-$S$ vintage is created, and the inherited one is offered at ψ to nobody. In every period, the input types in use, the prices of intermediate inputs and of agricultural goods, revenues, profits, and incomes are therefore the objects of Definition A.1 at the one-period model's decisions, which proves the claim. □

Remark: patents that no successor supersedes. For a type-S patent created at $\lambda_{t-1} \geq e^{-\delta}$ that no successor has superseded by period $t+n$, the argument of the proof of Proposition 5(i) applies with $\widehat{\lambda}_t = \lambda_{t-1}$, the innovation being non-drastic: the holder posts its original limit price, its price-adjusted quality relative to the frontier's is $\widehat{\lambda}_t/(1+\xi)^n$, and Southern farmers choose its vintage as long as this ratio is at or above $e^{-\delta}$, so the patent earns its creation margin on Southern input sales in those periods and nothing afterwards. When type-S innovation is unprofitable at some quality ratio at or above $e^{-\delta}$ under one-period patents, so that successors need not enter, a longer patent can therefore raise a holder's profit, but only by postponing the disclosure of its blueprint: farmers receive the improvement at marginal cost at expiry or supersession instead of one period after its creation, and the additional profit comes from delaying the knowledge spillover to households described in Section 3.3. The additional profit can induce an innovation that would not occur under one-period patents. With two-period patents, let $\bar{\pi}_{S,t}(\lambda)$ denote the discounted operating profits of a type-S patent created at λ that no successor supersedes: the first-period profit $\pi_{S,t}(\lambda)$ plus the discounted profit of the second period. If $\pi_{S,t}(\lambda) < c \leq \bar{\pi}_{S,t}(\lambda)$ and $\bar{\pi}_{S,t}(T(\lambda)) < c$, the firm that inherits $T(\lambda)$ after an innovation at λ does not enter, whatever the firms of later periods do, because even a patent that no successor supersedes would not cover the fixed cost there. The patent created at λ is then not superseded in its second period and earns $\bar{\pi}_{S,t}(\lambda) \geq c$, so type-$S$ innovation occurs at λ although one period's operating profit does not cover the fixed cost; by Proposition 5(ii), it is not followed by another in the next period. Numerical examples that satisfy Assumptions A.1–A.5 meet this condition. How often such innovations occur and what they imply for welfare are quantitative questions that I do not pursue.

B Proofs and Derivations for the Quantitative Model

This appendix contains the derivations and proofs for the quantitative model of Section 4. Section B.1 states the assumptions formally. Section B.2 solves the farmer's problem and aggregates over farmers. Section B.3 derives the prices of intermediate inputs, the price-adjusted quality ratio (44), and the properties of the ratio $\widetilde{\mathcal{Q}}$. Section B.4 derives profits and discusses the equilibrium innovation intensity. Section B.5 derives the law of motion (42). Section B.6 collects the equations that determine the equilibrium of a given period. Section B.7 proves Proposition 6, derives its counterpart with the relative market access M_t, and derives the corresponding expression for agricultural employment shares. Section B.8 considers two special cases: the quantitative model without trade costs, and the baseline model.

Throughout, as in Appendix A, $\psi = \alpha^2$, $\gamma = 1 - \alpha - \beta$, $\Gamma = 1 + \gamma(\varepsilon - 1)$, and the two agricultural varieties have equal weights, $\mu = 1/2$. I write $\mu_N := \mu$ and $\mu_S := 1 - \mu$ for the weights of the Northern and the Southern variety in the aggregator (14), so that $\mu_k = 1/2$ throughout, and I keep them explicit where a derivation does not depend on their value. I use the following symbols, which do not appear in the main text. As in Appendix A, $R_{j,t} := p^a_{j,t} Y^a_{j,t}$ denotes agricultural revenue in country j, $R^W_t := R_{N,t} + R_{S,t}$ world agricultural revenue, and $s_{S,t} := R_{S,t}/R^W_t$ the South's share of world agricultural revenue. Let

$$h^*_{j,t} := \left(h^\theta_{j,N,t} + h^\theta_{j,S,t}\right)^{1/\theta}$$

summarize the price-adjusted effective qualities of the two input types in country j. As $\theta \to \infty$, $h^*_{j,t}$ converges to the price-adjusted effective quality of the input type with the higher $h_{j,k,t}$, which is the object denoted by $h^*_{j,t}$ in Appendix A. The remaining symbols are introduced where they are first used.

B.1 Assumptions

Section 4 introduces the assumptions of the quantitative model where they are first used. This section collects them and states them formally. They extend Assumptions A.2–A.4 of the baseline model. The tie-breaking rule in Assumption A.1 is not needed, because a farmer is indifferent between the two input types with probability zero.

Assumption B.1.

(a) Suitability. *For each farmer $i \in [0,1]$ in country j, each input type k, and each period t, the suitability $v_{j,k,t}(i)$ is drawn independently across farmers, input types, and periods from the Fréchet distribution $\Pr\{v_{j,k,t}(i) \leq x\} = \exp(-x^{-\theta})$ for $x > 0$, with $\theta > \alpha/\gamma$. The share of farmers in country j whose draws have a given property equals the probability of that property.*

(b) One input type per farmer. *Each farmer uses one input type on all product lines.*

Part (a) implies that both input types are used in each country and that the adoption shares are smooth functions of the price-adjusted effective qualities (Lemma B.2). The restriction $\theta > \alpha/\gamma$ ensures that aggregate agricultural revenue is finite. Part (b) restricts the farmer's choice only when type-S innovation is drastic and some, but not all, product lines of type-S inputs are improved; otherwise it is implied by optimality (the end of Section B.3).

Assumption B.2.

(a) Disclosure and knowledge sharing at patent expiry. *When the patents granted in period t expire, the blueprints of the improved product lines become publicly available, and the knowledge embodied in them is shared within the input type. From period $t+1$ onward, the competitive fringe can produce every product line of type-k inputs at the base quality $q_{k,t}$ in (31) at marginal cost ψ. In period t, the competitive fringe can produce every product line of type-k inputs at the base quality $q_{k,t-1}$, and a product line that is not improved in period t is sold at price ψ.*

(b) Order of decisions within Stage 2. *Stage 2 consists of three steps. In Step 1, farmers choose an input type and hire labor, anticipating the prices of intermediate inputs. In Step 2, the patent holder of each improved product line posts a price, and the competitive fringe offers every product line at the base quality $q_{k,t-1}$ at marginal cost ψ. In Step 3, for each product line of the input type chosen in Step 1, farmers choose between the improved product line and the one offered by the competitive fringe and decide how much to purchase. A farmer who is indifferent between the two buys from the patent holder. Production, trade, and consumption then take place. Farmers have rational expectations, so the anticipated and the posted prices coincide in equilibrium.*

(c) Perceived demand. *When posting its price in Step 2, the patent holder of a product line takes as given agricultural output prices, household incomes, and the choices that farmers made in Step 1.*

Assumption B.2 applies Assumption A.2 to each product line. The only addition is the sharing of knowledge in part (a): at the beginning of each period, all product lines of an input type have the same base quality, so the base qualities of the two input types summarize the state of technology (Section 4.1.2).

Assumption B.3.

(a) Non-drastic frontier innovation. $1 + \xi < 1/\alpha$.

(b) All product lines of type-N inputs are improved. *In every period t, $\pi_{N,t} \geq \bar{c}_N$, whatever the innovation intensity of type-S inputs in that period.*

(c) Location of R&D. *All R&D firms are located in the North.*

(d) R&D firms and managers. *In each period, one R&D firm per input type holds the research opportunities on all product lines of that input type. Each product line is run by a manager, who observes the R&D fixed cost $c_{k,\ell,t}$ and improves the product line if and only if $\pi_{k,t} \geq c_{k,\ell,t}$, taking*

as given aggregate variables and the decisions of the other managers. The R&D firm earns the operating profits of the one-period patents on its improved product lines and exits at the end of the period. It holds no claim on research opportunities in later periods. Northern households own the R&D firms.

(e) R&D fixed costs. *For each input type k, the ratio $c_{k,\ell,t}/c_k$ is drawn independently across product lines, input types, and periods from a distribution with mean one and cumulative distribution function H. The function H is continuous, and it is strictly increasing on its support $[\underline{c}_k/c_k, \bar{c}_k/c_k]$, where $0 < \underline{c}_k < \bar{c}_k < \infty$ and the two ratios $\underline{c}_k/c_k$ and $\bar{c}_k/c_k$ are the same for both input types. The share of product lines of type-k inputs whose R&D fixed cost is at most x equals $H(x/c_k)$.*

(f) Imitators. $\zeta_S \in (0,1]$. *Imitators in the South receive the fraction $1-\zeta_S$ of the operating profit that the patent holder of each improved product line earns from sales to Southern farmers. Imitation affects neither the prices nor the quantities of intermediate inputs. Southern households own the imitators.*

Parts (a) and (c) are as in Assumption A.3. Part (b) is the assumption introduced in Section 4.2 to focus on type-S innovation, and it is the counterpart of Assumption A.3(b). Part (d) implies that the entry decision on each product line compares the operating profit of a single period with the R&D fixed cost of that product line, as in Assumption A.3(d). The same product lines are improved if, instead, each product line is open to the pool of potential R&D firms of Section 2.1.2: the first firm that enters a product line obtains the patent, and it enters if and only if $\pi_{k,t} \geq c_{k,\ell,t}$, because it takes aggregate variables as given, as a manager does. In part (e), the continuity of H implies that the innovation intensity (35) varies continuously with the operating profit, and $\underline{c}_k > 0$ implies that no product line is improved if the operating profit is sufficiently low. Part (f) treats imitation as a transfer from patent holders to Southern households: imitators do not produce, so they do not affect the prices paid by farmers.

Assumption B.4.

(a) Parameters and initial condition. *The parameters and the initial base qualities satisfy Assumption A.4(a). In addition, $Z_j > 0$ and $\tau_{w,j} \geq 0$ for $j \in \{N,S\}$, $f \geq 0$, $\bar{a} \geq 0$, and $\tau_g \geq 1$.*

(b) Positive world income net of R&D fixed costs. $A_N^n L_N + A_S^n L_S > (2+f)\, A_N^n/\eta$.

(c) Interior allocation. *In every period, $L_{j,t}^n > 0$ and $E_{j,t} > P_{j,t}^a \bar{a} L_j$ for $j \in \{N,S\}$.*

Part (b) is the counterpart of Assumption A.4(b): the R&D fixed costs paid in a period are at most $c_N + c_S = (2+f)\, A_N^n/\eta$, which is their value if all product lines of both input types are improved. The first condition in part (c) implies that the non-agricultural wage is $w_{j,t}^n = A_j^n$ in both countries. The second condition states that household income exceeds the cost of subsistence consumption, so that the household's problem has a solution with $C_{j,t}^a > \bar{a} L_j$ and $C_{j,t}^n > 0$. Part (c) restricts equilibrium objects. Appendix C.9 verifies it, together with Assumption B.3(b),

at the calibrated parameters. Finally, because households hold no assets, the household's problem is static: in each period, the household maximizes the consumption aggregate (36) subject to the budget constraint (37). Section 4 and this appendix use only this static problem.

B.2 Farmers: Input Demand, Input Type Choice, and Aggregation

This section solves the farmer's problem of Section 4.1.1 and aggregates over farmers. Fix a period t and a country j, and take as given the agricultural output price $p^a_{j,t}$, the agricultural wage $w^a_{j,t}$, and the prices of intermediate inputs. In equilibrium, farmers anticipate the posted prices (Assumption B.2(b)), so their choices coincide with the solution to the farmer's problem evaluated at these prices, as in Section A.2.

Expenditure across product lines. The first step is the allocation of a farmer's expenditure on intermediate inputs across the product lines of the input type that the farmer uses.

Lemma B.1. *Consider a farmer in country j who uses type-k inputs, has suitability v for them, and spends X on intermediate inputs. The farmer's expenditure on product line ℓ is*

$$p^x_{k,\ell,t}\, x_\ell = X \left[\frac{q_{k,\ell,t} / p^x_{k,\ell,t}}{\widetilde{q}_{k,t}} \right]^{\frac{\alpha}{1-\alpha}}, \tag{B.1}$$

where x_ℓ denotes the quantity of product line ℓ that the farmer purchases, and

$$\int_0^1 \left(v\, B_{j,k}\, q_{k,\ell,t}\, x_\ell\right)^\alpha \mathrm{d}\ell = \left(v\, h_{j,k,t}\, X\right)^\alpha. \tag{B.2}$$

Proof. The farmer chooses $\{x_\ell\}_{\ell\in[0,1]}$ to maximize $\int_0^1 (v\, B_{j,k}\, q_{k,\ell,t}\, x_\ell)^\alpha \,\mathrm{d}\ell$ subject to $\int_0^1 p^x_{k,\ell,t}\, x_\ell \,\mathrm{d}\ell = X$. The objective is strictly concave, so the first-order conditions characterize the solution. With $\Lambda > 0$ denoting the multiplier on the constraint, the first-order condition for x_ℓ is $\alpha\,(v\, B_{j,k}\, q_{k,\ell,t})^\alpha\, x_\ell^{\alpha-1} = \Lambda\, p^x_{k,\ell,t}$, which gives

$$p^x_{k,\ell,t}\, x_\ell = \left(\frac{\alpha}{\Lambda}\right)^{\frac{1}{1-\alpha}} \left(v\, B_{j,k}\right)^{\frac{\alpha}{1-\alpha}} \left(\frac{q_{k,\ell,t}}{p^x_{k,\ell,t}}\right)^{\frac{\alpha}{1-\alpha}}.$$

Expenditure on a product line is therefore proportional to $(q_{k,\ell,t}/p^x_{k,\ell,t})^{\alpha/(1-\alpha)}$. Integrating over product lines, and using the constraint and the definition of $\widetilde{q}_{k,t}$ in (29), gives (B.1). Substituting x_ℓ from (B.1) gives

$$\left(v\, B_{j,k}\, q_{k,\ell,t}\, x_\ell\right)^\alpha = \left(v\, B_{j,k}\, X\right)^\alpha \widetilde{q}_{k,t}^{\,-\frac{\alpha^2}{1-\alpha}} \left(\frac{q_{k,\ell,t}}{p^x_{k,\ell,t}}\right)^{\alpha+\frac{\alpha^2}{1-\alpha}},$$

where the last exponent equals $\alpha/(1-\alpha)$. By the definition of $\widetilde{q}_{k,t}$, the integral of the last factor over product lines is $\widetilde{q}_{k,t}^{\,\alpha/(1-\alpha)}$. The integral of the left-hand side is therefore $(v\, B_{j,k}\, X)^\alpha\, \widetilde{q}^{\,\alpha}_{k,t}$, because $\alpha/(1-\alpha) - \alpha^2/(1-\alpha) = \alpha$. This is (B.2), because $h_{j,k,t} = B_{j,k}\,\widetilde{q}_{k,t}$. □

By (B.2), a farmer who spends X on type-k inputs obtains $v\,h_{j,k,t}\,X$ effective units of intermediate inputs. The farmer is therefore in the same position as a farmer in the baseline model who buys a single input with the price-adjusted effective quality $v\,h_{j,k,t}$. This is the sense in which $\widetilde{q}_{k,t}$ measures the quality of type-k inputs per unit of expenditure.

Labor, expenditure on intermediate inputs, and revenue of a farmer. By Lemma B.1, farmer i in country j, if using type-k inputs, chooses labor L and expenditure on intermediate inputs X to maximize

$$p_{j,t}^{a}\,A_j^{a}\,Z_j^{\gamma}\,L^{\beta}\left(v_{j,k,t}(i)\,h_{j,k,t}\,X\right)^{\alpha}-w_{j,t}^{a}\,L-X.$$

The first-order conditions are

$$X=\alpha\,p_{j,t}^{a}\,Y_{j,t}^{a}(i)\qquad\text{and}\qquad w_{j,t}^{a}\,L_{j,t}^{a}(i)=\beta\,p_{j,t}^{a}\,Y_{j,t}^{a}(i),$$

which are the demands stated in Section 4.1.1. Substituting them into the production function and multiplying by $p_{j,t}^{a}$ gives $p_{j,t}^{a}Y_{j,t}^{a}(i)=p_{j,t}^{a}A_j^{a}Z_j^{\gamma}\left(\beta\,p_{j,t}^{a}Y_{j,t}^{a}(i)/w_{j,t}^{a}\right)^{\beta}\left(\alpha\,v_{j,k,t}(i)\,h_{j,k,t}\,p_{j,t}^{a}Y_{j,t}^{a}(i)\right)^{\alpha}$, and hence the revenue of farmer i:

$$p_{j,t}^{a}\,Y_{j,t}^{a}(i)=Z_j\left[\alpha^{\alpha}\beta^{\beta}A_j^{a}\left(w_{j,t}^{a}\right)^{-\beta}p_{j,t}^{a}\right]^{1/\gamma}\left(v_{j,k,t}(i)\,h_{j,k,t}\right)^{\alpha/\gamma}.\tag{B.3}$$

The farmer spends the share α of revenue on intermediate inputs and the share β on labor. The remaining share γ is the farmer's profit, which is the income from land.

Input type choice. By (B.3), the farmer's profit is strictly increasing in $v_{j,k,t}(i)\,h_{j,k,t}$. Farmer i therefore uses the input type with the highest $v_{j,k,t}(i)\,h_{j,k,t}$, as stated in Section 4.1.1.

Lemma B.2. *Suppose that Assumption B.1 holds. In country j in period t:*

(i) The share of farmers who use type-k inputs is $\omega_{j,k,t}=h_{j,k,t}^{\theta}/\left(h_{j,N,t}^{\theta}+h_{j,S,t}^{\theta}\right)$, which is (30).

(ii) For every $y>0$, the share of farmers who use type-k inputs and for whom $\max_{k'}v_{j,k',t}(i)\,h_{j,k',t}\le y$ is $\omega_{j,k,t}\exp[-(h_{j,t}^{*}/y)^{\theta}]$. Hence the distribution of $\max_{k'}v_{j,k',t}(i)\,h_{j,k',t}$ among the farmers who use type-k inputs is the same for both input types.

(iii) $\displaystyle\int_0^1\left(\max_k v_{j,k,t}(i)\,h_{j,k,t}\right)^{\alpha/\gamma}\mathrm{d}i=\left(h_{j,t}^{*}\right)^{\alpha/\gamma}\int_0^{\infty}u^{-\frac{\alpha}{\gamma\theta}}e^{-u}\,\mathrm{d}u<\infty.$

Proof. Fix an input type k, and let $o\neq k$ denote the other input type. I suppress the indices j and t. The product v_kh_k has the cumulative distribution function $\exp[-(h_k/z)^{\theta}]$ for $z>0$, with density $\theta\,h_k^{\theta}z^{-\theta-1}\exp[-(h_k/z)^{\theta}]$. The farmer uses type-$k$ inputs if $v_oh_o<v_kh_k$, and the two products are equal with probability zero. Because the draws are independent across input types, the probability that the farmer uses type-k inputs and that $v_kh_k\le y$ is

$$\begin{aligned}\int_0^{y}\exp\left[-(h_o/z)^{\theta}\right]\,\theta\,h_k^{\theta}z^{-\theta-1}\exp\left[-(h_k/z)^{\theta}\right]\mathrm{d}z&=\frac{h_k^{\theta}}{(h^{*})^{\theta}}\int_0^{y}\theta\,(h^{*})^{\theta}\,z^{-\theta-1}\exp\left[-(h^{*}/z)^{\theta}\right]\mathrm{d}z\\&=\omega_k\exp\left[-(h^{*}/y)^{\theta}\right],\end{aligned}$$

where the first equality uses $h_k^\theta + h_o^\theta = (h^*)^\theta$, and the second holds because the integrand is the derivative of $\exp[-(h^*/z)^\theta]$ with respect to z. By Assumption B.1(a), this probability equals the corresponding share of farmers, which gives part (ii). Letting $y \to \infty$ gives part (i).

Summing the expression in part (ii) over the two input types shows that $\max_k v_k h_k$ has the cumulative distribution function $\exp[-(h^*/y)^\theta]$. Hence $u := (h^*/\max_k v_k h_k)^\theta$ has the exponential distribution with mean one, and $(\max_k v_k h_k)^{\alpha/\gamma} = (h^*)^{\alpha/\gamma} u^{-\alpha/(\gamma\theta)}$. Taking the mean gives the expression in part (iii). The integral is finite if and only if $\alpha/(\gamma\theta) < 1$, which holds by Assumption B.1(a). □

Aggregation. Integrating (B.3) over farmers, using Lemma B.2(iii), and raising both sides to the power γ gives

$$R_{j,t}^\gamma = \Gamma_\theta \, \alpha^\alpha \beta^\beta A_j^a \, Z_j^\gamma \left(w_{j,t}^a\right)^{-\beta} p_{j,t}^a \left(h_{j,t}^*\right)^\alpha, \qquad \text{where} \qquad \Gamma_\theta := \left[\int_0^\infty u^{-\frac{\alpha}{\gamma\theta}} e^{-u} \, \mathrm{d}u\right]^\gamma. \tag{B.4}$$

Equation (B.4) is the counterpart of (A.1) in the baseline model. It differs in three respects. First, it includes the land endowment Z_j. Second, $h_{j,t}^*$ summarizes the price-adjusted effective qualities of both input types. Third, it includes the constant Γ_θ, which is common to the two countries and converges to one as $\theta \to \infty$. Four implications follow.

First, integrating the labor demand over farmers gives $w_{j,t}^a L_{j,t}^a = \beta R_{j,t}$. With $w_{j,t}^a = A_j^n/(1+\tau_{w,j})$ by (41) and $w_{j,t}^n = A_j^n$,

$$L_{j,t}^a = \frac{\beta\,(1+\tau_{w,j})\,R_{j,t}}{A_j^n} \qquad \text{and} \qquad \frac{Y_{j,t}^a}{L_{j,t}^a} = \frac{w_{j,t}^a}{\beta\, p_{j,t}^a} = \frac{A_j^n}{(1+\tau_{w,j})\,\beta\, p_{j,t}^a}, \tag{B.5}$$

where the second expression uses $Y_{j,t}^a = R_{j,t}/p_{j,t}^a$. As in (A.2), agricultural labor productivity equals the agricultural wage divided by $\beta p_{j,t}^a$.

Second, by (B.3), a farmer's revenue depends on the farmer's draws only through $\max_k v_{j,k,t}(i)\, h_{j,k,t}$. By Lemma B.2(ii), its distribution is the same among the farmers who use each input type. The farmers who use type-k inputs therefore earn the share $\omega_{j,k,t}$ of agricultural revenue in country j, and they account for the same share of agricultural labor and of the expenditure on intermediate inputs. Farmers in country j spend $\alpha\,\omega_{j,k,t}\,R_{j,t}$ on type-k inputs, so the adoption share is also the share of type-k inputs in the country's total expenditure on intermediate inputs. By (B.1), their expenditure on product line ℓ of type-k inputs is

$$\alpha\,\omega_{j,k,t}\,R_{j,t}\left[\frac{q_{k,\ell,t}/p_{k,\ell,t}^x}{\widetilde{q}_{k,t}}\right]^{\frac{\alpha}{1-\alpha}}. \tag{B.6}$$

Third, the profits of farmers are $\text{Profit}_{j,t}^a = \gamma R_{j,t}$.

Fourth, multiplying both sides of (B.4) by $R_{j,t}^{\alpha+\beta}/p_{j,t}^a$ and using $w_{j,t}^a L_{j,t}^a = \beta R_{j,t}$ gives

$$Y_{j,t}^a = \Gamma_\theta \, A_j^a \, Z_j^\gamma \left(L_{j,t}^a\right)^\beta \left(h_{j,t}^* \cdot \alpha R_{j,t}\right)^\alpha.$$

Aggregate agricultural output in country j therefore has the form of the production function (1) of the baseline model, where $h^*_{j,t}$ multiplies the country's total expenditure on intermediate inputs, $\alpha R_{j,t}$.

B.3 Prices and the Price-Adjusted Quality Ratio

Prices. The price of each product line is determined as in the baseline model.

Lemma B.3. *Suppose that Assumption B.2 holds. If product line ℓ of type-k inputs is improved in period t, the patent holder posts the price $p^x_{k,\ell,t} = \min\{\alpha,\ G_{k,t}\,\psi\}$. If it is not improved, it is sold at price ψ. The price of an improved product line depends neither on the size of its market nor on which country's farmers buy it.*

Proof. Consider Step 3 of Stage 2. A farmer in country j who chose type-k inputs and labor L in Step 1, with suitability v, buys x units of product line ℓ at price p' and base quality q. Because the production function (27) is additive across product lines, the farmer's payoff from product line ℓ is $p^a_{j,t} A^a_j Z^\gamma_j L^\beta (v\,B_{j,k}\,q\,x)^\alpha - p'x$, and it does not depend on the farmer's purchases of other product lines. Maximizing over x gives a quantity proportional to $(p')^{-1/(1-\alpha)}$ and the payoff

$$(1-\alpha)\,\alpha^{\frac{\alpha}{1-\alpha}}\left(p^a_{j,t} A^a_j Z^\gamma_j L^\beta\right)^{\frac{1}{1-\alpha}}\left(\frac{v\,B_{j,k}\,q}{p'}\right)^{\frac{\alpha}{1-\alpha}},$$

which is strictly increasing in q/p'. If product line ℓ is improved, the farmer can buy it at base quality $q_{k,\ell,t} = G_{k,t}\,q_{k,t-1}$ from the patent holder at price p, or at base quality $q_{k,t-1}$ from the competitive fringe at price ψ (Assumption B.2(a)). The farmer therefore buys from the patent holder if and only if $q_{k,\ell,t}/p \geq q_{k,t-1}/\psi$, that is, $p \leq G_{k,t}\,\psi$, as in (A.3). Under Assumption B.2(c), the patent holder thus faces a demand proportional to $p^{-1/(1-\alpha)}$ for $p \leq G_{k,t}\,\psi$ and zero demand for $p > G_{k,t}\,\psi$. This is the demand in the proof of Lemma A.1, and the remainder of that proof applies to the product line without change. □

Price-adjusted qualities. The following lemma derives the price-adjusted qualities of the two input types and the price-adjusted quality ratio $\widehat{\lambda}_t = \widetilde{q}_{S,t}/\widetilde{q}_{N,t}$.

Lemma B.4. *Suppose that Assumptions B.2 and B.3(a)–(b) hold. Fix a period t with inherited quality ratio $\lambda_{t-1} \in (0,\chi]$.*

(i) $\widetilde{q}_{N,t} = q_{N,t-1}/\psi$.

(ii) The price-adjusted quality of type-S inputs is

$$\widetilde{q}_{S,t} = \frac{q_{S,t-1}}{\psi}\left[m_{S,t}\max\left\{1,\ \alpha\,G_S(\lambda_{t-1})\right\}^{\frac{\alpha}{1-\alpha}} + 1 - m_{S,t}\right]^{\frac{1-\alpha}{\alpha}},$$

and the price-adjusted quality ratio $\widehat{\lambda}_t$ is given by (44).

(iii) $\lambda_{t-1} \le \widehat{\lambda}_t \le \chi$. *The ratio* $\widehat{\lambda}_t$ *is continuous in* λ_{t-1} *and* $m_{S,t}$. *It is strictly increasing in* $m_{S,t}$ *if* $\lambda_{t-1} < \lambda^D$, *and it does not depend on* $m_{S,t}$ *if* $\lambda_{t-1} \ge \lambda^D$.

(iv) Let $r_{k,t}$ *denote farmers' expenditure on an improved product line of type-k inputs relative to their average expenditure on a product line of type-k inputs. Then* $r_{N,t} = 1$ *and*

$$r_{S,t} = \left[\frac{\max\{\lambda_{t-1},\, \alpha(1+\xi)\,T(\lambda_{t-1})\}}{\widehat{\lambda}_t}\right]^{\frac{\alpha}{1-\alpha}} \ge 1,$$

with $r_{S,t} = 1$ *if* $\lambda_{t-1} \ge \lambda^D$. *Moreover,* $m_{S,t}\, r_{S,t} \le 1$.

Proof. (i) By Assumption B.3(a)–(b) and Lemma B.3, every product line of type-N inputs is improved and sold at price $(1+\xi)\,\psi$, with base quality $(1+\xi)\,q_{N,t-1}$. Hence $q_{N,\ell,t}/p^x_{N,\ell,t} = q_{N,t-1}/\psi$ for every ℓ, and (29) gives $\widetilde{q}_{N,t} = q_{N,t-1}/\psi$.

(ii) A product line of type-S inputs that is not improved has $q_{S,\ell,t}/p^x_{S,\ell,t} = q_{S,t-1}/\psi$. An improved product line has base quality $G_S(\lambda_{t-1})\,q_{S,t-1}$ and price $\min\{\alpha,\, G_S(\lambda_{t-1})\,\psi\}$. If innovation is non-drastic, its price is $G_S(\lambda_{t-1})\,\psi$ and $q_{S,\ell,t}/p^x_{S,\ell,t} = q_{S,t-1}/\psi$. If innovation is drastic, its price is α and $q_{S,\ell,t}/p^x_{S,\ell,t} = G_S(\lambda_{t-1})\,q_{S,t-1}/\alpha = \alpha\, G_S(\lambda_{t-1})\,q_{S,t-1}/\psi$, because $\psi = \alpha^2$. In both cases, $q_{S,\ell,t}/p^x_{S,\ell,t} = \max\{1,\, \alpha\, G_S(\lambda_{t-1})\}\,q_{S,t-1}/\psi$. Because the share $m_{S,t}$ of product lines is improved, (29) gives the expression for $\widetilde{q}_{S,t}$. Dividing by $\widetilde{q}_{N,t} = q_{N,t-1}/\psi$ gives

$$\widehat{\lambda}_t = \lambda_{t-1}\left[m_{S,t}\max\{1,\, \alpha\, G_S(\lambda_{t-1})\}^{\frac{\alpha}{1-\alpha}} + 1 - m_{S,t}\right]^{\frac{1-\alpha}{\alpha}}.$$

By Lemma 1, $\alpha\, G_S(\lambda_{t-1}) \ge 1$ if and only if $\lambda_{t-1} \le \lambda^{\mathrm{D}}$. If $\lambda_{t-1} > \lambda^{\mathrm{D}}$, the expression equals λ_{t-1}. If $\lambda_{t-1} \le \lambda^{\mathrm{D}}$, then $\lambda_{t-1}\,\alpha\, G_S(\lambda_{t-1}) = \alpha(1+\xi)\,T(\lambda_{t-1})$ by (18), which gives the second line of (44). The two lines coincide at $\lambda_{t-1} = \lambda^{\mathrm{D}}$, where $\alpha\, G_S(\lambda^{\mathrm{D}}) = 1$.

(iii) The term in square brackets is at least one, so $\widehat{\lambda}_t \ge \lambda_{t-1}$. It is at most $\max\{1,\, \alpha\, G_S(\lambda_{t-1})\}^{\alpha/(1-\alpha)}$, so $\widehat{\lambda}_t \le \max\{\lambda_{t-1},\, \alpha(1+\xi)\,T(\lambda_{t-1})\} \le \chi$, where the last inequality uses $\alpha(1+\xi)\,T(\lambda_{t-1}) \le \alpha(1+\xi)\,\chi < \chi$. Continuity follows from the continuity of G_S. If $\lambda_{t-1} < \lambda^{\mathrm{D}}$, then $\alpha\, G_S(\lambda_{t-1}) > 1$ and the term in square brackets is strictly increasing in $m_{S,t}$. If $\lambda_{t-1} \ge \lambda^{\mathrm{D}}$, it equals one.

(iv) By (B.1), the expenditure on product line ℓ relative to the average expenditure on a product line of type-k inputs is $[(q_{k,\ell,t}/p^x_{k,\ell,t})/\widetilde{q}_{k,t}]^{\alpha/(1-\alpha)}$. For type-$N$ inputs, this ratio is one by part (i). For an improved product line of type-S inputs, part (ii) gives $(q_{S,\ell,t}/p^x_{S,\ell,t})/\widetilde{q}_{S,t} = \max\{1,\, \alpha\, G_S(\lambda_{t-1})\}\,\lambda_{t-1}/\widehat{\lambda}_t$, which yields the expression for $r_{S,t}$. It is at least one because $\widehat{\lambda}_t \le \max\{\lambda_{t-1},\, \alpha(1+\xi)\,T(\lambda_{t-1})\}$ by the proof of part (iii). Finally, by part (ii), $m_{S,t}\, r_{S,t}$ equals the ratio of $m_{S,t}\max\{1,\alpha\, G_S(\lambda_{t-1})\}^{\alpha/(1-\alpha)}$ to $m_{S,t}\max\{1,\alpha\, G_S(\lambda_{t-1})\}^{\alpha/(1-\alpha)} + 1 - m_{S,t}$, which is at most one. □

When type-S innovation is non-drastic, every product line of type-S inputs offers the same quality per unit of expenditure, whether or not it is improved, so farmers spend the same

amount on each of them ($r_{S,t} = 1$) and the price-adjusted quality ratio equals the inherited quality ratio. When type-S innovation is drastic, an improved product line offers a higher quality per unit of expenditure, and farmers spend more on it than on a product line that is not improved ($r_{S,t} > 1$ if $m_{S,t} < 1$). The product $m_{S,t} r_{S,t}$ is the share of improved product lines in the expenditure on type-S inputs.

Adoption shares and the effective quality ratio. The next lemma expresses the adoption shares and the composites $h^*_{j,t}$ in terms of the price-adjusted quality ratio and establishes the properties of the ratio $\widetilde{\mathcal{Q}}$ in (46).

Lemma B.5. *Suppose that Assumptions B.1–B.3 hold. In period t:*

(i) $h^*_{N,t} = \dfrac{q_{N,t-1}}{\psi}\left[1 + (e^{-\delta}\widehat{\lambda}_t)^{\theta}\right]^{1/\theta}$ *and* $h^*_{S,t} = \dfrac{q_{N,t-1}}{\psi}\left[e^{-\delta\theta} + \widehat{\lambda}_t^{\theta}\right]^{1/\theta}$. *Hence* $h^*_{S,t}/h^*_{N,t} = \widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$.

(ii) The adoption shares of type-S inputs are

$$\omega_{S,S,t} = \frac{\widehat{\lambda}_t^{\theta}}{e^{-\delta\theta} + \widehat{\lambda}_t^{\theta}} \qquad \text{and} \qquad \omega_{N,S,t} = \frac{(e^{-\delta}\widehat{\lambda}_t)^{\theta}}{1 + (e^{-\delta}\widehat{\lambda}_t)^{\theta}},$$

and $\omega_{j,N,t} = 1 - \omega_{j,S,t}$. *Both are strictly increasing in* $\widehat{\lambda}_t$, *and* $0 < \omega_{N,S,t} < \omega_{S,S,t} < 1$.

(iii) The function $\widetilde{\mathcal{Q}}$ is strictly increasing, with elasticity

$$\frac{\mathrm{d}\ln\widetilde{\mathcal{Q}}(\widehat{\lambda})}{\mathrm{d}\ln\widehat{\lambda}} = \frac{\widehat{\lambda}^{\theta}}{e^{-\delta\theta} + \widehat{\lambda}^{\theta}} - \frac{(e^{-\delta}\widehat{\lambda})^{\theta}}{1 + (e^{-\delta}\widehat{\lambda})^{\theta}} > 0,$$

which equals $\omega_{S,S,t} - \omega_{N,S,t}$ *at* $\widehat{\lambda} = \widehat{\lambda}_t$. *Moreover,* $\widetilde{\mathcal{Q}}(\widehat{\lambda}) \to e^{-\delta}$ *as* $\widehat{\lambda} \to 0$ *and* $\widetilde{\mathcal{Q}}(\widehat{\lambda}) \to e^{\delta}$ *as* $\widehat{\lambda} \to \infty$.

(iv) For every $\widehat{\lambda} \in (0,1]$, $\widetilde{\mathcal{Q}}(\widehat{\lambda}) \geq \mathcal{Q}(\widehat{\lambda})$, *with equality if and only if* $\widehat{\lambda} = 1$, *where $\mathcal{Q}$ is the function defined in Proposition 1.*

(v) For every $\widehat{\lambda} \in (0,1]$, $\widetilde{\mathcal{Q}}(\widehat{\lambda}) \to \mathcal{Q}(\widehat{\lambda})$ *as* $\theta \to \infty$.

Proof. (i)–(ii) By (29) and (2), $h_{N,N,t} = \widetilde{q}_{N,t}$, $h_{N,S,t} = e^{-\delta}\widetilde{q}_{S,t} = e^{-\delta}\widehat{\lambda}_t\widetilde{q}_{N,t}$, $h_{S,N,t} = e^{-\delta}\widetilde{q}_{N,t}$, and $h_{S,S,t} = \widetilde{q}_{S,t} = \widehat{\lambda}_t\widetilde{q}_{N,t}$, where $\widetilde{q}_{N,t} = q_{N,t-1}/\psi$ by Lemma B.4(i). Substituting these expressions into the definition of $h^*_{j,t}$ gives part (i), and the ratio of the two composites is (46). Substituting them into (30) gives the adoption shares in part (ii). Writing $\omega_{S,S,t} = [1 + (e^{\delta}\widehat{\lambda}_t)^{-\theta}]^{-1}$ and $\omega_{N,S,t} = [1 + (e^{-\delta}\widehat{\lambda}_t)^{-\theta}]^{-1}$ shows that both are strictly increasing in $\widehat{\lambda}_t$ and lie strictly between zero and one, and that $\omega_{N,S,t} < \omega_{S,S,t}$ because $\delta > 0$.

(iii) By (46), $\ln\widetilde{\mathcal{Q}}(\widehat{\lambda}) = \theta^{-1}\{\ln[e^{-\delta\theta} + \widehat{\lambda}^{\theta}] - \ln[1 + (e^{-\delta}\widehat{\lambda})^{\theta}]\}$. Differentiating with respect to $\ln\widehat{\lambda}$ gives the elasticity, which is strictly positive by the inequality in part (ii). The two limits follow from (46).

(iv) If $\widehat{\lambda} \in [e^{-\delta}, 1]$, then $\mathcal{Q}(\widehat{\lambda}) = \widehat{\lambda}$, and $\widetilde{\mathcal{Q}}(\widehat{\lambda}) \geq \widehat{\lambda}$ is equivalent to $e^{-\delta\theta} + \widehat{\lambda}^{\theta} \geq \widehat{\lambda}^{\theta}[1 + e^{-\delta\theta}\widehat{\lambda}^{\theta}]$, that is, to $\widehat{\lambda}^{2\theta} \leq 1$, which holds with equality if and only if $\widehat{\lambda} = 1$. If $\widehat{\lambda} < e^{-\delta}$, then $\mathcal{Q}(\widehat{\lambda}) = e^{-\delta}$, and $\widetilde{\mathcal{Q}}(\widehat{\lambda}) > e^{-\delta}$ by part (iii).

(v) For $x, y > 0$, $(x^{\theta} + y^{\theta})^{1/\theta} \to \max\{x, y\}$ as $\theta \to \infty$. Hence $\widetilde{\mathcal{Q}}(\widehat{\lambda}) \to \max\{e^{-\delta}, \widehat{\lambda}\} / \max\{1, e^{-\delta}\widehat{\lambda}\}$. For $\widehat{\lambda} \leq 1$, the denominator equals one, and the limit is $e^{-\delta}$ if $\widehat{\lambda} < e^{-\delta}$ and $\widehat{\lambda}$ if $\widehat{\lambda} \geq e^{-\delta}$, which is $\mathcal{Q}(\widehat{\lambda})$. □

Two remarks are in order. First, by Lemma B.4(iii), $\widehat{\lambda}_t \leq \chi \leq 1$ in equilibrium, so parts (iv) and (v) apply at every price-adjusted quality ratio that arises in equilibrium. Second, $(x^{\theta} + y^{\theta})^{1/\theta} \geq \max\{x, y\}$ implies that each composite $h^*_{j,t}$ is at least as high as the price-adjusted effective quality of the better input type, which is its value in the baseline model. Part (iv) states that the difference is proportionally larger in the South than in the North.

One input type per farmer. Suppose that a farmer could choose the input type separately for each product line. On product line ℓ, the farmer would use the input type with the higher $v_{j,k,t}(i)\, B_{j,k}\, q_{k,\ell,t} / p^x_{k,\ell,t}$. When type-$S$ innovation is non-drastic, or when either all or none of the product lines of type-S inputs are improved, $q_{k,\ell,t} / p^x_{k,\ell,t} = \widetilde{q}_{k,t}$ on every product line of both input types (Lemma B.4). The comparison is then the same on every product line, the farmer uses one input type on all product lines, and Assumption B.1(b) does not restrict the farmer's choice. When type-S innovation is drastic and $0 < m_{S,t} < 1$, the improved product lines of type-S inputs offer a higher quality per unit of expenditure than the other product lines of type-S inputs. A farmer for whom type-S inputs are preferable on the improved product lines but not on the others would then use type-S inputs on the former and type-N inputs on the latter. Assumption B.1(b) rules this out. It reflects that the choice of an input type is made before the season begins and depends on seed systems, extension services, and learning that are specific to an input type, as discussed after Assumption A.2.

B.4 Profits and the Equilibrium Innovation Intensity

Operating profit of an improved product line. By (B.6) and Lemma B.4(iv), farmers in country j spend $r_{k,t} \cdot \alpha\, \omega_{j,k,t}\, R_{j,t}$ on an improved product line of type-k inputs. As in Section A.2, the operating profit from these sales is the margin rate $\nu(G_{k,t})$ times this expenditure. The patent holder keeps the entire operating profit from sales to Northern farmers and the fraction ζ_S of the operating profit from sales to Southern farmers (Assumption B.3(f)). The operating profit of the patent holder of an improved product line of type-k inputs is therefore

$$\pi_{k,t} = \nu(G_{k,t})\, r_{k,t} \left[\omega_{N,k,t} \cdot \alpha R_{N,t} + \zeta_S\, \omega_{S,k,t} \cdot \alpha R_{S,t}\right], \qquad \text{where} \qquad r_{k,t} = \left[\frac{q_{k,\ell,t} / p^x_{k,\ell,t}}{\widetilde{q}_{k,t}}\right]^{\frac{\alpha}{1-\alpha}} \tag{B.7}$$

for an improved product line ℓ. If type-k innovation is non-drastic, $r_{k,t} = 1$ and (B.7) is (32). If type-S innovation is drastic, the margin rate is $1 - \alpha$ and $r_{S,t} = [\alpha(1 + \xi)\, T(\lambda_{t-1}) / \widehat{\lambda}_t]^{\alpha/(1-\alpha)}$,

which is the additional factor mentioned in Section 4.1.2. It equals $[\alpha\, G_S(\lambda_{t-1})]^{\alpha/(1-\alpha)}$ if $m_{S,t} = 0$, it is strictly decreasing in $m_{S,t}$ if $\lambda_{t-1} < \lambda^{\mathrm{D}}$, and it equals one if $m_{S,t} = 1$: the fewer product lines are improved, the larger is the share of farmers' expenditure that each improved product line attracts.

Profits of imitators, net profits, and R&D fixed costs. Imitators receive the fraction $1 - \zeta_S$ of the operating profit from sales to Southern farmers on each improved product line. Their profits are therefore

$$\mathrm{Profit}^x_{S,t} = (1 - \zeta_S) \sum_{k \in \{N,S\}} m_{k,t}\, \nu(G_{k,t})\, r_{k,t}\, \omega_{S,k,t} \cdot \alpha R_{S,t}. \tag{B.8}$$

By Assumption B.3(d)–(e), the improved product lines of type-k inputs are those with $c_{k,\ell,t} \leq \pi_{k,t}$, and their share is $m_{k,t} = H(\pi_{k,t}/c_k)$, which is (35). Let H^{-1} denote the inverse of H on its support. The product lines with the lowest R&D fixed costs are improved, so the R&D fixed costs paid for type-k inputs are $c_k \int_0^{m_{k,t}} H^{-1}(u)\,\mathrm{d}u$, and the net profit of the R&D firm of type-k inputs is

$$\Pi_{k,t} = m_{k,t}\, \pi_{k,t} - c_k \int_0^{m_{k,t}} H^{-1}(u)\,\mathrm{d}u \;\geq\; 0. \tag{B.9}$$

The inequality holds because $c_k H^{-1}(u) \leq \pi_{k,t}$ for every $u \leq m_{k,t}$. Researchers are paid the R&D fixed costs, so

$$\begin{aligned} w^r_{N,t} L^r_{N,t} &= \sum_{k \in \{N,S\}} c_k \int_0^{m_{k,t}} H^{-1}(u)\,\mathrm{d}u, \\ w^r_{N,t} L^r_{N,t} + \sum_{k \in \{N,S\}} \Pi_{k,t} &= \sum_{k \in \{N,S\}} m_{k,t}\, \pi_{k,t}. \end{aligned} \tag{B.10}$$

The wages of researchers and the net profits of the R&D firms add up to the operating profits of the improved product lines. Because the mean of $c_{k,\ell,t}/c_k$ is one, the R&D fixed costs paid for type-k inputs are at most c_k.

Equilibrium innovation intensity. Fix a period t and the inherited base qualities. Under Assumption B.3(b), $m_{N,t} = 1$. For $m \in [0,1]$, let $\pi_{S,t}(m)$ denote the operating profit (B.7) of an improved product line of type-S inputs in the equilibrium of period t in which the share m of the product lines of type-S inputs is improved (Section B.6). By (35), m is an *equilibrium innovation intensity* of type-S inputs if and only if

$$m = H\!\left(\frac{\pi_{S,t}(m)}{c_S}\right). \tag{B.11}$$

The operating profit depends on m through three channels. First, household incomes, and therefore agricultural revenues, depend on the R&D fixed costs paid and on the profits earned on the improved product lines. Second, if type-S innovation is drastic, the price-adjusted quality ratio (44) and hence the adoption shares of type-S inputs increase with m. Third, if type-S

innovation is drastic, the relative expenditure $r_{S,t}$ decreases with m. Because of the first two channels, the operating profit of an improved product line can increase with the share of product lines that are improved, and more than one innovation intensity may satisfy (B.11). As stated in Section 4.2, I select the lowest one.

Lemma B.6. *Suppose that Assumptions B.1–B.4 and Condition B.1 hold. Then an equilibrium innovation intensity of type-S inputs exists in every period, and the set of equilibrium innovation intensities has a smallest element.*

Proof. By Condition B.1, the agricultural revenues $R_{N,t}$ and $R_{S,t}$ are continuous in m. By Lemmas B.4 and B.5, the price-adjusted quality ratio, the adoption shares, and $r_{S,t}$ are continuous in m. Hence $\pi_{S,t}(m)$ in (B.7) is continuous on $[0,1]$. Because H is continuous (Assumption B.3(e)), the function $m \mapsto H(\pi_{S,t}(m)/c_S) - m$ is continuous on $[0,1]$. It is nonnegative at $m = 0$ and nonpositive at $m = 1$, so it has a zero by the intermediate value theorem. The set of its zeros is closed and bounded, so it has a smallest element. □

Three remarks are in order. First, Condition B.1 holds without trade costs (Proposition B.1). Second, Lemma B.6 does not establish that the equilibrium innovation intensity is unique. Appendix C.9 verifies uniqueness at the calibrated parameters. Third, if $\pi_{S,t}(0) \leq \underline{c}_S$, then $m = 0$ satisfies (B.11), so no product line of type-S inputs is improved. Here, $\pi_{S,t}(0)$ is the operating profit that the manager of a product line anticipates if this product line alone is improved. The following lemma shows that this is the case when the inherited quality ratio is sufficiently low.

Lemma B.7. *Suppose that Assumptions B.1–B.4 and Condition B.1 hold.*

(i) In the equilibrium of period t, for any innovation intensities, $R_t^W \leq (A_N^n L_N + A_S^n L_S)/\alpha^2$.

(ii) For every inherited quality ratio $\lambda_{t-1} \in (0, \chi]$,

$$\pi_{S,t}(0) \leq \frac{1-\alpha}{\alpha} \max\left\{1, \frac{\alpha(1+\xi)\chi}{\lambda_{t-1}}\right\}^{\frac{\alpha}{1-\alpha}} \left(e^{\delta}\lambda_{t-1}\right)^{\theta} \left(A_N^n L_N + A_S^n L_S\right).$$

(iii) There exists $\underline{\lambda} > 0$, which depends only on α, ξ, χ, δ, θ, $A_N^n L_N + A_S^n L_S$, and $\underline{c}_S$, such that $\pi_{S,t}(0) < \underline{c}_S$ whenever $\lambda_{t-1} < \underline{\lambda}$. If $\lambda_{t-1} < \underline{\lambda}$, then $m_{S,s} = 0$ and $\lambda_s = \lambda_{s-1}/(1+\xi)$ in every period $s \geq t$, and the quality ratio converges to zero.

Proof. (i) By Lemma B.3, the price of every product line is at most α: it is $\min\{\alpha, G_{k,t}\psi\}$ or $\psi = \alpha^2$. Farmers spend αR_t^W on intermediate inputs (Section B.2), so they purchase at least R_t^W units. Producing them requires at least $\psi R_t^W = \alpha^2 R_t^W$ units of the non-agricultural good. By the market clearing condition for the non-agricultural good and $C_{j,t}^n \geq 0$, this quantity cannot exceed $Y_{N,t}^n + Y_{S,t}^n \leq A_N^n L_N + A_S^n L_S$.

(ii) If $m_{S,t} = 0$, then $\widehat{\lambda}_t = \lambda_{t-1}$ and $r_{S,t} = \max\{1, \alpha G_S(\lambda_{t-1})\}^{\alpha/(1-\alpha)}$ by Lemma B.4, where $\alpha G_S(\lambda_{t-1}) = \alpha(1+\xi)T(\lambda_{t-1})/\lambda_{t-1} \leq \alpha(1+\xi)\chi/\lambda_{t-1}$. Moreover, $\nu(G_S(\lambda_{t-1})) \leq 1-\alpha$ by (9),

$\zeta_S \leq 1$, and $\omega_{N,S,t} < \omega_{S,S,t} \leq (e^\delta \lambda_{t-1})^\theta$ by Lemma B.5(ii). Substituting these bounds into (B.7) gives $\pi_{S,t}(0) \leq (1-\alpha)\max\{1,\ \alpha(1+\xi)\chi/\lambda_{t-1}\}^{\alpha/(1-\alpha)}(e^\delta \lambda_{t-1})^\theta \cdot \alpha R_t^W$, and part (i) gives the bound.

(iii) For $\lambda_{t-1} \leq \alpha(1+\xi)\chi$, the bound in part (ii) equals a positive constant times $\lambda_{t-1}^{\theta-\alpha/(1-\alpha)}$. The exponent is strictly positive, because $\theta > \alpha/\gamma$ by Assumption B.1(a) and $\gamma < 1-\alpha$. The bound is therefore strictly increasing in λ_{t-1} on this interval and converges to zero as $\lambda_{t-1} \to 0$. Let $\underline{\lambda}$ be the smaller of $\alpha(1+\xi)\chi$ and the quality ratio at which the bound equals $\underline{c}_S$. Then $\pi_{S,t}(0) < \underline{c}_S$ for every $\lambda_{t-1} < \underline{\lambda}$, and $\underline{\lambda}$ does not depend on t, because $\underline{c}_S$ is constant over time. For such λ_{t-1}, $H(\pi_{S,t}(0)/c_S) = 0$, so $m = 0$ satisfies (B.11) and is the smallest equilibrium innovation intensity. By (42), $\lambda_t = \lambda_{t-1}/(1+\xi) < \underline{\lambda}$, and the same argument applies in period $t+1$ and in every later period. □

Lemma B.7 is the counterpart of the baseline result that type-S innovation does not occur when the inherited quality ratio is below $e^{-\delta}$ (Corollary A.1). In the baseline model, type-S inputs have no market at such quality ratios. In the quantitative model, some farmers use type-S inputs at every quality ratio, but their share is at most $(e^\delta \lambda_{t-1})^\theta$ and therefore vanishes as the quality ratio falls. A lower quality ratio also raises the potential for catch-up, which raises farmers' expenditure on a first improved product line relative to the other product lines of type-S inputs, by a factor that is at most proportional to $\lambda_{t-1}^{-\alpha/(1-\alpha)}$. Because $\theta > \alpha/(1-\alpha)$, the first effect dominates. Because the lowest R&D fixed cost is strictly positive, the operating profit then falls short of the R&D fixed cost of every product line.

B.5 Law of Motion of the Quality Ratio

Lemma B.8. *Suppose that Assumptions B.3(a)–(b) and B.4(a) hold, and let $\lambda_{t-1} \in (0, \chi]$.*

(i) The base quality of type-k inputs evolves according to $q_{k,t} = (G_{k,t})^{m_{k,t}} q_{k,t-1}$, and the quality ratio evolves according to (42).

(ii) $\lambda_{t-1}/(1+\xi) \leq \lambda_t \leq T(\lambda_{t-1}) \leq \chi$, and λ_t is strictly increasing in $m_{S,t}$. Hence the quality ratio remains in $(0, \chi]$ in every period.

(iii) $\lambda_t > \lambda_{t-1}$ if and only if $m_{S,t} > \overline{m}(\lambda_{t-1})$, and $\lambda_t = \lambda_{t-1}$ if and only if $m_{S,t} = \overline{m}(\lambda_{t-1})$. The function $\overline{m}$ in (43) is continuous and strictly increasing on $(0, \chi]$, with $\overline{m}(\lambda) \to 0$ as $\lambda \to 0$ and $\overline{m}(\chi) = 1$.

(iv) If $m_{S,t} = 1$ in every period, then $\chi - \lambda_t = (1-a)^{t+1}(\chi - \lambda_{-1})$, so the quality ratio converges to χ. If $m_{S,t} = 0$ in every period, then $\lambda_t = \lambda_{-1}/(1+\xi)^{t+1}$, so the quality ratio converges to zero.

Proof. (i) In period t, the share $m_{k,t}$ of the product lines of type-k inputs has the base quality $G_{k,t} q_{k,t-1}$, and the remaining product lines have the base quality $q_{k,t-1}$ (Section 4.1.2). By (31),

$\ln q_{k,t} = m_{k,t} \ln(G_{k,t}\, q_{k,t-1}) + (1 - m_{k,t}) \ln q_{k,t-1} = \ln q_{k,t-1} + m_{k,t} \ln G_{k,t}$. With $G_{N,t} = 1 + \xi$, $m_{N,t} = 1$ by Assumption B.3(b), and $G_{S,t} = (1 + \xi)\, T(\lambda_{t-1}) / \lambda_{t-1}$ by (18),

$$\lambda_t = \frac{q_{S,t}}{q_{N,t}} = \frac{\lambda_{t-1}}{1+\xi} \left[\frac{(1+\xi)\, T(\lambda_{t-1})}{\lambda_{t-1}} \right]^{m_{S,t}} = T(\lambda_{t-1})^{m_{S,t}} \left(\frac{\lambda_{t-1}}{1+\xi} \right)^{1-m_{S,t}}.$$

(ii) Because $T(\lambda_{t-1}) - \lambda_{t-1} = a\,(\chi - \lambda_{t-1}) \geq 0$, it follows that $T(\lambda_{t-1}) \geq \lambda_{t-1} > \lambda_{t-1}/(1+\xi)$ and $T(\lambda_{t-1}) \leq \chi$. By part (i), λ_t is a weighted geometric mean of $T(\lambda_{t-1})$ and $\lambda_{t-1}/(1+\xi)$ with the weight $m_{S,t}$ on the larger of the two, which gives the bounds and the monotonicity. Because $\lambda_{-1} \in (0, \chi]$ by Assumption B.4(a), the quality ratio remains in $(0, \chi]$ by induction. Hence $q_{S,t-1} \leq \chi\, q_{N,t-1}$ in every period, which is the condition under which catch-up takes the form in Section 4.1.2.

(iii) By part (i), $\ln \lambda_t - \ln \lambda_{t-1} = m_{S,t} \ln G_S(\lambda_{t-1}) - \ln(1+\xi)$. If $\lambda_{t-1} < \chi$, then $G_S(\lambda_{t-1}) > 1 + \xi$ by Lemma 1, so $\ln G_S(\lambda_{t-1}) > 0$ and the claims follow from the definition of $\overline{m}$. If $\lambda_{t-1} = \chi$, then $G_S(\chi) = 1 + \xi$ and $\overline{m}(\chi) = 1$: the quality ratio remains at χ if $m_{S,t} = 1$ and falls otherwise. Because G_S is continuous and strictly decreasing, with $G_S(\lambda) \to \infty$ as $\lambda \to 0$, the function $\overline{m}$ is continuous and strictly increasing, with $\overline{m}(\lambda) \to 0$ as $\lambda \to 0$.

(iv) If $m_{S,t} = 1$, then $\lambda_t = T(\lambda_{t-1})$ and $\chi - \lambda_t = (1-a)\,(\chi - \lambda_{t-1})$. If $m_{S,t} = 0$, then $\lambda_t = \lambda_{t-1}/(1+\xi)$. Iterating from λ_{-1} gives the two expressions. □

B.6 The Equilibrium of Period t

This section collects the equations that determine the equilibrium of period t for given inherited base qualities $(q_{N,t-1}, q_{S,t-1})$ and given innovation intensities $(m_{N,t}, m_{S,t}) = (1, m)$. The equilibrium of period t in Definition 3 is the equilibrium of this section evaluated at the equilibrium innovation intensity of Section B.4.

Households. The household in country j maximizes the consumption aggregate (36) subject to the budget constraint (37). It pays the delivered price $\tau_{j,k}\, p^a_{k,t}$ per unit of the agricultural variety produced in country k. Consider first the composition of agricultural consumption. Minimizing the expenditure $\sum_k \tau_{j,k}\, p^a_{k,t}\, c^a_{j,k,t}$ for a given $C^a_{j,t}$ in (14) gives $c^a_{j,k,t} = \mu_k\, (\tau_{j,k}\, p^a_{k,t} / P^a_{j,t})^{-\varepsilon}\, C^a_{j,t}$, where $P^a_{j,t}$ is the price index in (38). Hence the expenditure on agricultural goods is $E^a_{j,t} = P^a_{j,t}\, C^a_{j,t}$, and the expenditure of country j on the variety produced in country k is

$$\tau_{j,k}\, p^a_{k,t}\, c^a_{j,k,t} = \mu_k \left(\frac{\tau_{j,k}\, p^a_{k,t}}{P^a_{j,t}} \right)^{1-\varepsilon} E^a_{j,t}. \tag{B.12}$$

Consider next the division of income between agricultural goods and the non-agricultural good. The household chooses $C^a_{j,t}$ to maximize $(C^a_{j,t} - \overline{a} L_j)^{\phi}\, (E_{j,t} - P^a_{j,t} C^a_{j,t})^{1-\phi}$. The first-order condition is $\phi\, (E_{j,t} - P^a_{j,t} C^a_{j,t}) = (1-\phi)\, P^a_{j,t}\, (C^a_{j,t} - \overline{a} L_j)$, which gives the expenditure on agricultural goods in (39):

$$E^a_{j,t} = \phi\, E_{j,t} + (1-\phi)\, P^a_{j,t}\, \overline{a}\, L_j.$$

Under Assumption B.4(c), $C^a_{j,t} > \bar{a} L_j$ and $C^n_{j,t} = E_{j,t} - E^a_{j,t} > 0$.

Household incomes. By (41), $w^n_{j,t} = A^n_j$, $w^r_{N,t} = A^n_N$, and the labor market clearing conditions, labor income in country j is $A^n_j L_j - (A^n_j - w^a_{j,t}) L^a_{j,t}$. Because $A^n_j - w^a_{j,t} = \tau_{w,j} w^a_{j,t}$ and $w^a_{j,t} L^a_{j,t} = \beta R_{j,t}$, it equals $A^n_j L_j - \beta \tau_{w,j} R_{j,t}$. Each agricultural worker earns $\tau_{w,j} w^a_{j,t}$ less than a non-agricultural worker, so labor income falls short of $A^n_j L_j$ by this amount times agricultural employment. With $\text{Profit}^a_{j,t} = \gamma R_{j,t}$, household incomes (40) become

$$\begin{aligned} E_{N,t} &= A^n_N L_N + (\gamma - \beta \tau_{w,N}) R_{N,t} + \sum_{k \in \{N,S\}} \Pi_{k,t}, \\ E_{S,t} &= A^n_S L_S + (\gamma - \beta \tau_{w,S}) R_{S,t} + \text{Profit}^x_{S,t}, \end{aligned} \tag{B.13}$$

where $\Pi_{k,t}$ and $\text{Profit}^x_{S,t}$ are given by (B.9) and (B.8).

The system of equations. The equilibrium of period t is computed in three steps. In the first step, the price-adjusted quality ratio $\widehat{\lambda}_t$, the composites $h^*_{j,t}$, the adoption shares $\omega_{j,k,t}$, the margin rates, and $r_{S,t}$ follow from λ_{t-1}, $q_{N,t-1}$, and m (Lemmas 1, B.4, and B.5). None of them depends on agricultural revenues. In the second step, the two agricultural revenues $(R_{N,t}, R_{S,t})$ solve the following system. By (B.4) and (41), the agricultural output prices are

$$p^a_{j,t} = \frac{R^\gamma_{j,t}}{\Gamma_\theta \, \alpha^\alpha \beta^\beta A^a_j Z^\gamma_j \left[A^n_j / (1 + \tau_{w,j})\right]^{-\beta} \left(h^*_{j,t}\right)^\alpha}, \qquad j \in \{N,S\}. \tag{B.14}$$

The price indices $P^a_{j,t}$ follow from (38), and the expenditures on agricultural goods $E^a_{j,t}$ follow from (39), with the household incomes in (B.13). Finally, by the market clearing condition for the agricultural good produced in country k and (B.12),

$$R_{k,t} = \sum_{j \in \{N,S\}} \tau_{j,k} \, p^a_{k,t} \, c^a_{j,k,t} = \mu_k \left(p^a_{k,t}\right)^{1-\varepsilon} \sum_{j \in \{N,S\}} \tau^{1-\varepsilon}_{j,k} E^a_{j,t} \left(P^a_{j,t}\right)^{\varepsilon - 1}, \qquad k \in \{N,S\}. \tag{B.15}$$

The first equality holds because the quantity shipped from country k to country j is $\tau_{j,k} \, c^a_{j,k,t}$, so the farmers in country k receive the entire expenditure of both countries on their variety. After substituting (B.14), (38), (39), and (B.13), the two equations in (B.15) contain only the two unknowns $R_{N,t}$ and $R_{S,t}$. In the third step, the remaining objects follow: agricultural employment from (B.5), the number of researchers from (B.10), non-agricultural employment from the labor market clearing conditions, agricultural output from $Y^a_{j,t} = R_{j,t} / p^a_{j,t}$, and consumption from the household's demands.

Condition B.1. *For all inherited base qualities with $\lambda_{t-1} \in (0, \chi]$ and every $m \in [0,1]$, the system (B.14)–(B.15) has a unique solution $(R_{N,t}, R_{S,t})$ with strictly positive revenues, and this solution is continuous in m.*

Without trade costs, Condition B.1 holds (Proposition B.1). With trade costs, the system has no closed-form solution, and I verify Condition B.1 numerically at the calibrated parameters (Appendix C.9). The following lemma shows that the system (B.14)–(B.15) implies market clearing for the non-agricultural good, and it relates world agricultural revenue to world income.

Lemma B.9. *Suppose that Assumptions B.1–B.4 hold and that* $(R_{N,t}, R_{S,t})$ *solves* (B.14)–(B.15).

(i) World agricultural revenue equals world expenditure on agricultural goods: $R_t^W = E_{N,t}^a + E_{S,t}^a$.

(ii) The market for the non-agricultural good clears.

(iii) World agricultural revenue satisfies

$$\widetilde{\kappa}_t R_t^W = \phi \left[A_N^n L_N + A_S^n L_S - w_{N,t}^r L_{N,t}^r \right] + (1-\phi)\,\bar{a} \sum_{j \in \{N,S\}} P_{j,t}^a L_j, \tag{B.16}$$

where

$$\begin{aligned} \widetilde{\kappa}_t \coloneqq 1 - \phi\gamma + \phi\beta \left[\tau_{w,N} (1 - s_{S,t}) + \tau_{w,S}\, s_{S,t} \right] \\ - \phi\alpha \sum_{k \in \{N,S\}} m_{k,t}\, \nu(G_{k,t})\, r_{k,t} \left[\omega_{N,k,t} (1 - s_{S,t}) + \omega_{S,k,t}\, s_{S,t} \right] > 0. \end{aligned} \tag{B.17}$$

Proof. (i) By the definition of $P_{j,t}^a$ in (38), $\sum_k \mu_k (\tau_{j,k}\, p_{k,t}^a / P_{j,t}^a)^{1-\varepsilon} = 1$. Summing (B.15) over k therefore gives $R_t^W = E_{N,t}^a + E_{S,t}^a$.

(ii) Farmers spend αR_t^W on intermediate inputs. A product line that is not improved is sold at marginal cost, so the expenditure on it equals the quantity of the non-agricultural good used to produce it. The expenditure on an improved product line equals this quantity plus the operating profit, which accrues to the patent holder and to imitators. Let Ψ_t denote the quantity of the non-agricultural good used to produce intermediate inputs, which is the third term on the left-hand side of the market clearing condition for the non-agricultural good in Section 4.1.3. Then

$$\alpha R_t^W = \Psi_t + \sum_{k \in \{N,S\}} m_{k,t}\, \pi_{k,t} + \text{Profit}_{S,t}^x.$$

Summing the household incomes (40) over the two countries, and using $w_{j,t}^a L_{j,t}^a = \beta R_{j,t}$, $w_{j,t}^n L_{j,t}^n = Y_{j,t}^n$, $\text{Profit}_{j,t}^a = \gamma R_{j,t}$, and (B.10), gives

$$\begin{aligned} E_{N,t} + E_{S,t} &= (\beta + \gamma)\, R_t^W + Y_{N,t}^n + Y_{S,t}^n + \sum_{k \in \{N,S\}} m_{k,t}\, \pi_{k,t} + \text{Profit}_{S,t}^x \\ &= R_t^W - \Psi_t + Y_{N,t}^n + Y_{S,t}^n, \end{aligned}$$

where the second equality uses the previous display and $\alpha + \beta + \gamma = 1$. World income thus equals the value added of all sectors. By the budget constraints (37) and part (i), $C_{N,t}^n + C_{S,t}^n = E_{N,t} + E_{S,t} - R_t^W = Y_{N,t}^n + Y_{S,t}^n - \Psi_t$, which is the market clearing condition for the non-agricultural good.

(iii) Summing (B.13) over the two countries and using (B.10) gives

$$\begin{aligned} E_{N,t} + E_{S,t} = A_N^n L_N + A_S^n L_S - w_{N,t}^r L_{N,t}^r + \gamma R_t^W - \beta \sum_{j \in \{N,S\}} \tau_{w,j}\, R_{j,t} \\ + \sum_{k \in \{N,S\}} m_{k,t}\, \pi_{k,t} + \text{Profit}_{S,t}^x. \end{aligned}$$

By (B.7) and (B.8), the last two terms add up to $\alpha \sum_k m_{k,t} \nu(G_{k,t}) r_{k,t} [\omega_{N,k,t} R_{N,t} + \omega_{S,k,t} R_{S,t}]$: the profits of patent holders and of imitators add up to the operating profits on all improved product lines. By part (i) and (39), $R_t^W = \phi (E_{N,t} + E_{S,t}) + (1 - \phi) \bar{a} \sum_j P_{j,t}^a L_j$. Substituting the expression for $E_{N,t} + E_{S,t}$, writing $R_{N,t} = (1 - s_{S,t}) R_t^W$ and $R_{S,t} = s_{S,t} R_t^W$, and collecting the terms in R_t^W gives (B.16).

It remains to show that $\widetilde{\kappa}_t > 0$. The term with the labor wedges is nonnegative. Every margin rate is at most $1 - \alpha$, $m_{k,t} r_{k,t} \leq 1$ by Lemma B.4(iv), and $\sum_k [\omega_{N,k,t} (1 - s_{S,t}) + \omega_{S,k,t} s_{S,t}] = 1$. Hence $\widetilde{\kappa}_t \geq 1 - \phi\gamma - \phi\alpha(1 - \alpha) > 1 - \phi(\gamma + \alpha) > 0$. □

Equation (B.16) is the counterpart of (A.5), and $\widetilde{\kappa}_t$ is the counterpart of κ_t in (A.6). There are two differences. First, the labor wedges raise $\widetilde{\kappa}_t$ and thus lower world agricultural revenue for a given right-hand side. With a labor wedge, higher agricultural revenue raises agricultural employment, and each agricultural worker earns less than a non-agricultural worker, so household income rises by less with agricultural revenue. Second, subsistence consumption adds a term that depends on the price indices. For this reason, and because, with trade costs, the share $s_{S,t}$ depends on the relative market access defined in Section B.7, (B.16) is not a closed-form solution in general. It is one without trade costs and subsistence consumption (Proposition B.1).

Lemma B.10. *Suppose that Assumptions B.1–B.4 hold. Fix the innovation intensities and a constant $g > 0$. The revenues $(R_{N,t}, R_{S,t})$ solve the system (B.14)–(B.15) at the inherited base qualities $(g q_{N,t-1}, g q_{S,t-1})$ and subsistence consumption $\bar{a}$ if and only if they solve it at the inherited base qualities $(q_{N,t-1}, q_{S,t-1})$ and subsistence consumption $g^{-\alpha} \bar{a}$. The two equilibria have the same adoption shares, agricultural revenues, household incomes, profits, and employment, and the agricultural output prices in the first are $g^{-\alpha}$ times those in the second.*

Proof. The price-adjusted quality ratio, the adoption shares, the margin rates, and $r_{S,t}$ depend on the inherited base qualities only through λ_{t-1}, which is the same in the two cases. By Lemma B.5(i), $h_{j,t}^*$ is proportional to $q_{N,t-1}$. For given revenues, (B.14) therefore implies that the agricultural output prices, and hence the price indices in (38), are $g^{-\alpha}$ times as high in the first case as in the second. The cost of subsistence consumption $P_{j,t}^a \bar{a} L_j$ is then the same in the two cases, because the first has the price index multiplied by $g^{-\alpha}$ and the second has subsistence consumption multiplied by $g^{-\alpha}$. Hence (B.13) and the expenditures on agricultural goods in (39) are the same. In (B.15), the factor $g^{-\alpha(1-\varepsilon)}$ from $(p_{k,t}^a)^{1-\varepsilon}$ and the factor $g^{-\alpha(\varepsilon-1)}$ from $(P_{j,t}^a)^{\varepsilon-1}$ cancel. The two systems therefore have the same solutions. Profits and employment depend on the solution only through revenues. □

Lemma B.10 implies that the equilibrium of period t, including the operating profit $\pi_{S,t}(m)$ and thus the equilibrium innovation intensity, depends on the inherited base qualities only through the quality ratio λ_{t-1} and the term $\bar{a} q_{N,t-1}^{-\alpha}$, as stated in Section 4.2. Without subsistence consumption, it depends only on λ_{t-1}, as in the baseline model. Because $q_{N,t-1}$ grows by the factor $1 + \xi$ per period (Lemma B.8(i)), the term $\bar{a} q_{N,t-1}^{-\alpha}$ declines by the factor $(1 + \xi)^{-\alpha}$ per period and converges to zero: subsistence consumption matters less as the frontier quality rises.

B.7 Proof of Proposition 6, Market Access, and Agricultural Employment Shares

This section proves Proposition 6, which rests on the supply side of the model alone. It then brings in the demand for agricultural goods: it derives the counterpart of (22) with the relative market access M_t, establishes the properties of M_t, discusses the two limiting cases of Section 4.3.1, and derives the ratio of agricultural employment shares. Each step states the equation of the previous sections on which it relies.

Proof of Proposition 6. Step 1: Agricultural labor productivity and the agricultural output price. By (B.5), $Y^a_{j,t}/L^a_{j,t} = w^a_{j,t}/(\beta\, p^a_{j,t})$. Dividing this expression for the South by the one for the North gives

$$\frac{Y^a_{S,t}/L^a_{S,t}}{Y^a_{N,t}/L^a_{N,t}} = \frac{w^a_{S,t}}{w^a_{N,t}} \cdot \frac{p^a_{N,t}}{p^a_{S,t}}, \qquad \text{where} \qquad \frac{w^a_{S,t}}{w^a_{N,t}} = \frac{A^n_S}{A^n_N} \cdot \frac{1+\tau_{w,N}}{1+\tau_{w,S}} \tag{B.18}$$

by (41) and $w^n_{j,t} = A^n_j$. Agricultural labor productivity differences equal the ratio of agricultural wages divided by the relative price of the Southern agricultural good. The ratio of agricultural wages consists of parameters. It therefore remains to determine the relative price $p^a_{S,t}/p^a_{N,t}$.

Step 2: Supply. Dividing (B.4) for the South by (B.4) for the North, the constants $\Gamma_\theta\, \alpha^\alpha \beta^\beta$ cancel, and

$$\left(\frac{R_{S,t}}{R_{N,t}}\right)^\gamma = \left(\frac{A^a_S}{A^a_N}\right)\left(\frac{Z_S}{Z_N}\right)^\gamma \left(\frac{w^a_{S,t}}{w^a_{N,t}}\right)^{-\beta} \left(\frac{h^*_{S,t}}{h^*_{N,t}}\right)^\alpha \frac{p^a_{S,t}}{p^a_{N,t}}.$$

By Lemma B.5(i), $h^*_{S,t}/h^*_{N,t} = \widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$. Let

$$\widetilde{\mathcal{F}}_t := \left(\frac{A^a_S}{A^a_N}\right)\left(\frac{Z_S}{Z_N}\right)^\gamma \left(\frac{A^n_S}{A^n_N} \cdot \frac{1+\tau_{w,N}}{1+\tau_{w,S}}\right)^{-\beta} \widetilde{\mathcal{Q}}(\widehat{\lambda}_t)^\alpha, \tag{B.19}$$

which is the counterpart of $\mathcal{F}_t$ in (A.4). Then

$$\left(\frac{R_{S,t}}{R_{N,t}}\right)^\gamma = \widetilde{\mathcal{F}}_t \frac{p^a_{S,t}}{p^a_{N,t}}. \tag{B.20}$$

For a given relative price, the country with stronger agricultural fundamentals, a lower agricultural wage, or a higher effective quality of intermediate inputs earns a higher agricultural revenue.

Step 3: Relative agricultural employment. Solving (B.20) for the relative price and substituting it into (B.18) gives

$$\frac{Y^a_{S,t}/L^a_{S,t}}{Y^a_{N,t}/L^a_{N,t}} = \left(\frac{A^a_S}{A^a_N}\right)\left(\frac{Z_S}{Z_N}\right)^\gamma \left(\frac{A^n_S}{A^n_N} \cdot \frac{1+\tau_{w,N}}{1+\tau_{w,S}}\right)^{\alpha+\gamma} \widetilde{\mathcal{Q}}(\widehat{\lambda}_t)^\alpha \left(\frac{R_{S,t}}{R_{N,t}}\right)^{-\gamma}, \tag{B.21}$$

where the exponent of the ratio of agricultural wages is $1-\beta = \alpha+\gamma$. By (B.5), $R_{j,t} = w^a_{j,t} L^a_{j,t}/\beta$, so $R_{S,t}/R_{N,t} = (w^a_{S,t}/w^a_{N,t})\,(L^a_{S,t}/L^a_{N,t})$. Substituting this ratio into (B.21) lowers the exponent of the ratio of agricultural wages to α and gives (45). □

The proof uses only (B.4), (B.5), and Lemma B.5(i): the production function and the farmers' demands for labor and intermediate inputs. It does not use the demand for agricultural goods. Proposition 6 therefore holds for any innovation intensities, whether type-S innovation is drastic or non-drastic, and for any trade cost and any level of subsistence consumption. For a given relative agricultural revenue, the elasticity of agricultural productivity differences with respect to $\widetilde{\mathcal{Q}}$ in (B.21) is α, the share of intermediate inputs in production. A higher relative agricultural revenue in the South means higher relative agricultural employment in the South, by (B.5), and lowers the South's relative agricultural labor productivity with the elasticity γ, the share of land.

Relative agricultural revenue and market access. Relative agricultural revenue, and with it relative agricultural employment, is determined by the demand for agricultural goods. Let

$$D_{j,t} \coloneqq E^a_{j,t} \left(P^a_{j,t}\right)^{\varepsilon-1} \qquad \text{and} \qquad M_{k,t} \coloneqq \sum_{j\in\{N,S\}} \tau_{j,k}^{1-\varepsilon} D_{j,t}. \tag{B.22}$$

The term $D_{j,t}$ measures the demand for agricultural goods in country j: it is high when the country's expenditure on agricultural goods is high and when agricultural goods are expensive there, so that each variety faces little competition. The term $M_{k,t}$ is the *market access* of farmers in country k: the sum of the demand in the two countries, where the demand in the other country is discounted by the trade cost. With $\tau_{N,N} = \tau_{S,S} = 1$ and $\tau_{N,S} = \tau_{S,N} = \tau_g$, $M_{N,t} = D_{N,t} + \tau_g^{1-\varepsilon} D_{S,t}$ and $M_{S,t} = \tau_g^{1-\varepsilon} D_{N,t} + D_{S,t}$. Let

$$M_t \coloneqq \frac{M_{S,t}}{M_{N,t}} = \frac{\tau_g^{1-\varepsilon} E^a_{N,t} \left(P^a_{N,t}\right)^{\varepsilon-1} + E^a_{S,t} \left(P^a_{S,t}\right)^{\varepsilon-1}}{E^a_{N,t} \left(P^a_{N,t}\right)^{\varepsilon-1} + \tau_g^{1-\varepsilon} E^a_{S,t} \left(P^a_{S,t}\right)^{\varepsilon-1}} \tag{B.23}$$

denote the South's market access relative to the North's.

Corollary B.1. *In equilibrium, agricultural labor productivity differences between the South and the North in period t are given by*

$$\frac{Y^a_{S,t}/L^a_{S,t}}{Y^a_{N,t}/L^a_{N,t}} = \left[\left(\frac{A^a_S}{A^a_N}\right) \cdot \left(\frac{Z_S}{Z_N}\right)^{\gamma} \cdot \left(\frac{A^n_S}{A^n_N} \cdot \frac{1+\tau_{w,N}}{1+\tau_{w,S}}\right)^{\alpha+\gamma\varepsilon} \cdot \widetilde{\mathcal{Q}}\left(\widehat{\lambda}_t\right)^{\alpha} \cdot M_t^{-\gamma} \right]^{1/\Gamma}, \tag{B.24}$$

where M_t is the relative market access in (B.23) and $\Gamma = 1 + \gamma(\varepsilon - 1)$, as in Proposition 1.

Proof. Step 1: Demand. By (B.15) with equal weights, $\mu_k = 1/2$, and the definitions in (B.22), the agricultural revenue of country k is

$$R_{k,t} = \frac{1}{2} \left(p^a_{k,t}\right)^{1-\varepsilon} M_{k,t}.$$

Because $\varepsilon > 1$, a lower price raises the revenue of country k for a given market access, and better market access raises it for a given price. Dividing this expression for the South by the one for

the North gives

$$\frac{R_{S,t}}{R_{N,t}} = \left(\frac{p^a_{S,t}}{p^a_{N,t}}\right)^{1-\varepsilon} M_t. \tag{B.25}$$

Without trade costs, $M_{N,t} = M_{S,t}$ and (B.25) is the relative demand in the proof of Proposition A.1(iv).

Step 2: Relative price. Raising (B.25) to the power γ and equating it with (B.20) gives $(p^a_{S,t}/p^a_{N,t})^{\gamma(1-\varepsilon)} M_t^{\gamma} = \widetilde{\mathcal{F}}_t (p^a_{S,t}/p^a_{N,t})$. Because $\gamma(1-\varepsilon) - 1 = -\Gamma$, this is $(p^a_{S,t}/p^a_{N,t})^{-\Gamma} = \widetilde{\mathcal{F}}_t M_t^{-\gamma}$, and hence

$$\frac{p^a_{S,t}}{p^a_{N,t}} = \left(\widetilde{\mathcal{F}}_t M_t^{-\gamma}\right)^{-1/\Gamma} \qquad \text{and} \qquad \frac{R_{S,t}}{R_{N,t}} = \widetilde{\mathcal{F}}_t^{(\varepsilon-1)/\Gamma} M_t^{1/\Gamma}, \tag{B.26}$$

where the second expression follows from substituting the first into (B.25) and using $1 - \gamma(\varepsilon - 1)/\Gamma = 1/\Gamma$.

Step 3: Agricultural labor productivity differences. Substituting the relative price in (B.26) into (B.18) gives

$$\frac{Y^a_{S,t}/L^a_{S,t}}{Y^a_{N,t}/L^a_{N,t}} = \frac{w^a_{S,t}}{w^a_{N,t}} \left(\widetilde{\mathcal{F}}_t M_t^{-\gamma}\right)^{1/\Gamma} = \left[\left(\frac{w^a_{S,t}}{w^a_{N,t}}\right)^{\Gamma} \widetilde{\mathcal{F}}_t M_t^{-\gamma}\right]^{1/\Gamma}.$$

By (B.19), the ratio of agricultural wages enters with the exponent $\Gamma - \beta = 1 + \gamma\varepsilon - \gamma - \beta = \alpha + \gamma\varepsilon$. Substituting (B.19) and the ratio of agricultural wages in (B.18) gives (B.24). □

Equation (B.24) is the counterpart of (22) in the baseline model, with the same exponents. It results from substituting the relative agricultural revenue in (B.26) into (B.21), and it holds for any innovation intensities and whether type-S innovation is drastic or non-drastic, because the innovation intensities and the prices of intermediate inputs affect agricultural productivity differences only through $\widehat{\lambda}_t$ and M_t. Unlike the other terms of (B.24), the relative market access is an equilibrium object. Farmers with better market access face a higher demand for their agricultural good. They earn higher agricultural revenue and employ more labor on a fixed amount of land, which lowers agricultural labor productivity. This is why M_t enters (B.24) with a negative exponent.

Properties of the relative market access. The next lemma collects the properties of M_t that the rest of this section uses.

Lemma B.11. *In the equilibrium of period t:*

(i) $M_t = \dfrac{D_{S,t}/D_{N,t} + \tau_g^{1-\varepsilon}}{\tau_g^{1-\varepsilon} D_{S,t}/D_{N,t} + 1}$. *If $\tau_g = 1$, then $M_t = 1$.*

(ii) If $\tau_g > 1$, M_t is strictly increasing in $D_{S,t}/D_{N,t}$, it lies between one and $D_{S,t}/D_{N,t}$, and $\tau_g^{1-\varepsilon} < M_t < \tau_g^{\varepsilon-1}$.

(iii) The ratio of the price indices $P^a_{S,t}/P^a_{N,t}$ is a function of the relative price $p^a_{S,t}/p^a_{N,t}$ alone. It is equal to one if $\tau_g = 1$ and strictly increasing in $p^a_{S,t}/p^a_{N,t}$ if $\tau_g > 1$.

(iv) For a given household income $E_{j,t}$, the expenditure on agricultural goods $E^a_{j,t}$ does not depend on the price index $P^a_{j,t}$ if $\bar{a} = 0$, and it is strictly increasing in $P^a_{j,t}$ if $\bar{a} > 0$.

Proof. Let $\vartheta := \tau_g^{1-\varepsilon} \in (0,1]$, with $\vartheta = 1$ if and only if $\tau_g = 1$, and let $d := D_{S,t}/D_{N,t}$. (i) Dividing the numerator and the denominator of $M_t = (D_{S,t} + \vartheta D_{N,t})/(\vartheta D_{S,t} + D_{N,t})$ by $D_{N,t}$ gives $M_t = (d+\vartheta)/(\vartheta d + 1)$, which equals one if $\vartheta = 1$. (ii) The derivative of $(d+\vartheta)/(\vartheta d+1)$ with respect to d is $(1-\vartheta^2)/(\vartheta d+1)^2 > 0$. Moreover, $M_t - 1 = (d-1)(1-\vartheta)/(\vartheta d+1)$ has the sign of $d-1$, and $M_t - d = \vartheta(1-d^2)/(\vartheta d+1)$ has the sign of $1-d$, so M_t lies between one and d. As d increases from zero to infinity, M_t increases from ϑ to $1/\vartheta$. (iii) By (38) with equal weights, $\mu_k = 1/2$, and with $u := (p^a_{S,t}/p^a_{N,t})^{1-\varepsilon}$,

$$\left(\frac{P^a_{S,t}}{P^a_{N,t}}\right)^{1-\varepsilon} = \frac{(p^a_{S,t})^{1-\varepsilon} + \vartheta\,(p^a_{N,t})^{1-\varepsilon}}{\vartheta\,(p^a_{S,t})^{1-\varepsilon} + (p^a_{N,t})^{1-\varepsilon}} = \frac{u+\vartheta}{\vartheta u + 1},$$

which equals one if $\vartheta = 1$ and is strictly increasing in u if $\vartheta < 1$, as in part (ii). Because $\varepsilon > 1$, u is strictly decreasing in $p^a_{S,t}/p^a_{N,t}$, and $P^a_{S,t}/P^a_{N,t}$ is strictly decreasing in $(P^a_{S,t}/P^a_{N,t})^{1-\varepsilon}$. The composition is strictly increasing. (iv) This follows from (39). □

Two limiting cases. Without trade costs ($\tau_g = 1$), $M_t = 1$ by Lemma B.11(i), and (B.24) becomes

$$\frac{Y^a_{S,t}/L^a_{S,t}}{Y^a_{N,t}/L^a_{N,t}} = \left[\left(\frac{A^a_S}{A^a_N}\right)\left(\frac{Z_S}{Z_N}\right)^{\gamma}\left(\frac{A^n_S}{A^n_N}\cdot\frac{1+\tau_{w,N}}{1+\tau_{w,S}}\right)^{\alpha+\gamma\varepsilon}\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)^{\alpha}\right]^{1/\Gamma},$$

which has the form of (22) and is a closed-form expression in parameters and the price-adjusted quality ratio. The elasticity of agricultural productivity differences with respect to $\widetilde{\mathcal{Q}}$ is α/Γ. It is smaller than the elasticity α in (B.21) because relative agricultural revenue responds to $\widetilde{\mathcal{Q}}$: by (B.26) with $M_t = 1$, $R_{S,t}/R_{N,t} = \widetilde{\mathcal{F}}_t^{(\varepsilon-1)/\Gamma}$, so an improvement in type-S inputs, which raises $\widetilde{\mathcal{Q}}$, raises the agricultural revenue and employment of the South relative to the North.

Consider next the economy without trade in agricultural goods, in which each household consumes only the domestic variety. This is the limiting case $\tau_g \to \infty$, in which $\tau_g^{1-\varepsilon}$ is replaced by zero. The price index in (38) is then $P^a_{j,t} = \mu_j^{1/(1-\varepsilon)} p^a_{j,t}$, $M_{k,t} = D_{k,t}$, and (B.15) becomes $R_{k,t} = \mu_k (p^a_{k,t})^{1-\varepsilon} E^a_{k,t} (P^a_{k,t})^{\varepsilon-1} = E^a_{k,t}$: the agricultural revenue of each country equals its own expenditure on agricultural goods. Equation (B.21) then gives

$$\frac{Y^a_{S,t}/L^a_{S,t}}{Y^a_{N,t}/L^a_{N,t}} = \left(\frac{A^a_S}{A^a_N}\right)\left(\frac{Z_S}{Z_N}\right)^{\gamma}\left(\frac{A^n_S}{A^n_N}\cdot\frac{1+\tau_{w,N}}{1+\tau_{w,S}}\right)^{\alpha+\gamma}\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)^{\alpha}\left(\frac{E^a_{S,t}}{E^a_{N,t}}\right)^{-\gamma}. \tag{B.27}$$

For a given relative expenditure on agricultural goods, the elasticity of agricultural productivity differences with respect to $\widetilde{\mathcal{Q}}$ is α. Without trade, an improvement in type-S inputs does not

allow Southern farmers to sell more abroad, so Southern agricultural revenue and employment change only to the extent that Southern expenditure on agricultural goods changes.

Market access and the quality ratio. With trade costs, M_t depends on the quality ratio through the ratio $D_{S,t}/D_{N,t} = (E^a_{S,t}/E^a_{N,t})(P^a_{S,t}/P^a_{N,t})^{\varepsilon-1}$. Suppose that type-$S$ inputs improve, so that $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$ and $\widetilde{\mathcal{F}}_t$ rise. For a given M_t, the relative price $p^a_{S,t}/p^a_{N,t}$ falls by (B.26): the Southern agricultural good becomes cheaper relative to the Northern one. Two channels then affect M_t. First, by Lemma B.11(iii), the ratio of the price indices $P^a_{S,t}/P^a_{N,t}$ falls if $\tau_g > 1$. The reason is that, with trade costs, the Southern agricultural good accounts for a larger share of the expenditure on agricultural goods in the South than in the North: by (B.12), these shares are $\mu_S\,(p^a_{S,t}/P^a_{S,t})^{1-\varepsilon}$ and $\mu_S\,(\tau_g\,p^a_{S,t}/P^a_{N,t})^{1-\varepsilon}$, and the first exceeds the second if and only if $(P^a_{S,t}/P^a_{N,t})^{1-\varepsilon} < \tau_g^{\varepsilon-1}$, which holds for $\tau_g > 1$ because the ratio $(u+\vartheta)/(\vartheta u+1)$ in the proof of Lemma B.11(iii) is below $1/\vartheta = \tau_g^{\varepsilon-1}$. For given expenditures on agricultural goods, a lower $P^a_{S,t}/P^a_{N,t}$ lowers $D_{S,t}/D_{N,t}$ and therefore M_t by Lemma B.11(ii). Second, by Lemma B.11(iv), with subsistence consumption a lower price index reduces the expenditure on agricultural goods for a given household income. The response of M_t also depends on the changes in household incomes, for which there is no closed-form expression. If M_t falls, the relative agricultural revenue $R_{S,t}/R_{N,t} = \widetilde{\mathcal{F}}_t^{(\varepsilon-1)/\Gamma}\,M_t^{1/\Gamma}$ rises by less than without trade costs, and the South's relative agricultural labor productivity (B.24) rises by more. Section 6 computes the responses of relative agricultural employment and of agricultural productivity differences at the calibrated parameters.

Agricultural employment shares. The same steps give the ratio of agricultural employment shares between the South and the North, which is the counterpart of Corollary A.2.

Corollary B.2. *Suppose that Assumptions B.1–B.4 hold. In equilibrium, the ratio of agricultural employment shares between the South and the North in period t is given by*

$$\frac{L^a_{S,t}/L_S}{L^a_{N,t}/L_N} = \frac{L_N}{L_S}\left(\frac{A^n_S}{A^n_N}\cdot\frac{1+\tau_{w,N}}{1+\tau_{w,S}}\right)^{-\left[1+\frac{\beta(\varepsilon-1)}{\Gamma}\right]}\left[\left(\frac{A^a_S}{A^a_N}\right)\left(\frac{Z_S}{Z_N}\right)^{\gamma}\right]^{\frac{\varepsilon-1}{\Gamma}}\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)^{\frac{\alpha(\varepsilon-1)}{\Gamma}}\,M_t^{1/\Gamma}. \tag{B.28}$$

Proof. By (B.5), $L^a_{j,t}/L_j = \beta R_{j,t}/(w^a_{j,t}L_j)$. By (B.26),

$$\frac{L^a_{S,t}/L_S}{L^a_{N,t}/L_N} = \frac{L_N}{L_S}\left(\frac{w^a_{S,t}}{w^a_{N,t}}\right)^{-1}\frac{R_{S,t}}{R_{N,t}} = \frac{L_N}{L_S}\left(\frac{w^a_{S,t}}{w^a_{N,t}}\right)^{-1}\widetilde{\mathcal{F}}_t^{(\varepsilon-1)/\Gamma}\,M_t^{1/\Gamma}.$$

Substituting (B.19) and the ratio of agricultural wages in (B.18) gives (B.28). □

As in (A.8), a higher agricultural wage in the South relative to the North lowers the Southern agricultural employment share relative to the Northern one, and stronger agricultural fundamentals and a higher ratio $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t)$ raise it. In addition, better market access of Southern farmers relative to Northern farmers raises it. Without trade costs, $M_t = 1$, so the ratio is a function of parameters and the effective quality ratio alone, as in (A.8).

B.8 Special Cases: No Trade Costs and the Baseline Model

No trade costs. Without trade costs, the equilibrium of period t for given innovation intensities can be characterized as in Proposition A.1.

Proposition B.1. *Suppose that Assumptions B.1–B.4 hold and that $\tau_g = 1$. Fix a period t, the inherited base qualities with $\lambda_{t-1} \in (0, \chi]$, and the innovation intensities $(m_{N,t}, m_{S,t}) = (1, m)$ with $m \in [0,1]$.*

(i) Relative price and agricultural revenue. $M_t = 1$, $P^a_{N,t} = P^a_{S,t}$, *and*

$$\frac{p^a_{S,t}}{p^a_{N,t}} = \widetilde{\mathcal{F}}_t^{-1/\Gamma}, \qquad \frac{R_{S,t}}{R_{N,t}} = \widetilde{\mathcal{F}}_t^{(\varepsilon-1)/\Gamma}, \qquad \text{and} \qquad s_{S,t} = \left[1 + \widetilde{\mathcal{F}}_t^{-(\varepsilon-1)/\Gamma}\right]^{-1},$$

where $\widetilde{\mathcal{F}}_t$ is given by (B.19).

(ii) World agricultural revenue without subsistence consumption. *If $\bar{a} = 0$, then*

$$R^W_t = \frac{\phi\left[A^n_N L_N + A^n_S L_S - w^r_{N,t} L^r_{N,t}\right]}{\widetilde{\kappa}_t},$$

where $\widetilde{\kappa}_t$ is given by (B.17) with $s_{S,t}$ from part (i), and $w^r_{N,t} L^r_{N,t}$ is given by (B.10).

(iii) Existence and uniqueness. *The system (B.14)–(B.15) has a unique solution with strictly positive revenues, and this solution is continuous in m. Hence Condition B.1 holds.*

Proof. (i) If $\tau_g = 1$, then $M_t = 1$ by Lemma B.11(i), and the two price indices in (38) coincide. The three expressions follow from (B.26) with $M_t = 1$ and the definition of $s_{S,t}$.

(ii) This is (B.16) with $\bar{a} = 0$. By part (i), $s_{S,t}$ does not depend on the level of world agricultural revenue, so neither does $\widetilde{\kappa}_t$.

(iii) By part (i), every solution has the share $s_{S,t}$ of part (i), which is uniquely determined by $\widetilde{\mathcal{F}}_t$. It remains to show that R^W_t is uniquely determined. Let P^a_t denote the common price index. By (B.14), $p^a_{N,t}$ and $p^a_{S,t}$ are proportional to $((1 - s_{S,t})\, R^W_t)^\gamma$ and $(s_{S,t}\, R^W_t)^\gamma$, so $P^a_t = b_t\, (R^W_t)^\gamma$, where $b_t > 0$ does not depend on R^W_t. By (B.16), R^W_t solves

$$\widetilde{\kappa}_t\, R^W_t - (1-\phi)\,\bar{a}\,(L_N + L_S)\, b_t \left(R^W_t\right)^\gamma = \phi\left[A^n_N L_N + A^n_S L_S - w^r_{N,t} L^r_{N,t}\right].$$

The right-hand side is strictly positive by Assumption B.4(b), because the R&D fixed costs paid are at most $c_N + c_S$. Because $\widetilde{\kappa}_t > 0$ and $\gamma \in (0,1)$, the left-hand side is a convex function of R^W_t that equals zero at $R^W_t = 0$ and tends to infinity as $R^W_t \to \infty$. If $\bar{a} = 0$, it is linear and strictly increasing. If $\bar{a} > 0$, it is strictly convex, so it is nonpositive up to its unique strictly positive zero and strictly increasing beyond it. In both cases, the equation has exactly one strictly positive solution, at which the left-hand side is strictly increasing. Conversely, the revenues $R_{N,t} = (1 - s_{S,t})\, R^W_t$ and $R_{S,t} = s_{S,t}\, R^W_t$ constructed from this solution solve the system. The share $s_{S,t}$, $\widetilde{\kappa}_t$, b_t, and $w^r_{N,t} L^r_{N,t}$ are continuous in m by Lemmas B.4 and B.5 and (B.10), so the solution is continuous in m. □

Proposition B.1 extends Proposition A.1(iv)–(v) to the quantitative model without trade costs. With subsistence consumption, world agricultural revenue solves one equation in one unknown, which has a unique solution. With trade costs, the relative price and the level of agricultural revenue must be solved jointly, because M_t depends on the expenditures on agricultural goods and on the price indices.

The baseline model as a special case. The baseline model of Section 2 is the quantitative model with a unit land endowment ($Z_j = 1$), full enforcement of intellectual property rights ($\zeta_S = 1$), no additional cost of improving type-S inputs ($f = 0$), no subsistence consumption ($\overline{a} = 0$), no trade costs ($\tau_g = 1$), and no labor wedges ($\tau_{w,j} = 0$), in the limit in which the dispersion of suitability across farmers and the dispersion of R&D fixed costs across product lines vanish. I consider the two limits in turn.

As $\theta \to \infty$, the dispersion of suitability vanishes. By Lemma B.5, $\omega_{N,S,t} \to 0$, $\omega_{S,S,t} \to 0$ if $\widehat{\lambda}_t < e^{-\delta}$, $\omega_{S,S,t} \to 1$ if $\widehat{\lambda}_t > e^{-\delta}$, and $\widetilde{\mathcal{Q}}(\widehat{\lambda}_t) \to \mathcal{Q}(\widehat{\lambda}_t)$: farmers' input choices and the ratio of the effective intermediate input quality converge to those of Proposition A.1(ii)–(iii). Moreover, $\Gamma_\theta \to 1$, and $h^*_{j,t}$ converges to the price-adjusted effective quality of the input type in use, so (B.4) becomes (A.1) with $Z_j = 1$ and $w^a_{j,t} = A^n_j$. With these restrictions, $\widetilde{\mathcal{F}}_t$ in (B.19) becomes $\mathcal{F}_t$ in (A.4), (45) holds with $\mathcal{Q}$ in place of $\widetilde{\mathcal{Q}}$, and Corollary B.1 with $M_t = 1$ becomes Proposition 1.

If all product lines of an input type have the same R&D fixed cost c, all managers make the same decision, so the innovation intensity is either zero or one. With $m_{S,t} \in \{0,1\}$, the price-adjusted quality ratio (44) becomes (21), the law of motion (42) becomes (17), $r_{S,t} = 1$ if $m_{S,t} = 1$, and the R&D fixed costs paid are $c\,(1 + m_{S,t})$. The operating profit (B.7) then becomes (8), household incomes (40) become (16), $\widetilde{\kappa}_t$ in (B.17) becomes κ_t in (A.6), and Proposition B.1(ii) becomes (A.5).

One difference concerns the entry decision. In the baseline model, one R&D firm decides whether to improve type-S inputs, and it compares the fixed cost of R&D with the operating profit that it earns if it innovates, which corresponds to $\pi_{S,t}(1)$. In the quantitative model, the manager of each product line takes the decisions of the other managers as given. If $\pi_{S,t}(0) < c \leq \pi_{S,t}(1)$, both $m_{S,t} = 0$ and $m_{S,t} = 1$ are consistent with the managers' decisions when all product lines have the same R&D fixed cost. The entry condition (12) of the baseline model then corresponds to the higher of the two innovation intensities, whereas Section 4.2 selects the lowest one. The two coincide whenever the equilibrium innovation intensity is unique.