Conceptual framework

The debate about the causes of structural transformation (defined as the shift of productive factors out of agriculture and into non-agricultural production) traces back to the early inception of the discipline that would become development economics in the post-war period. More specifically, this debate has substantially centered around identifying the effect of positive (productivity-increasing) shocks to the agricultural sector on the shift of factors across sectors; though more recently, a smaller literature has also analyzed the effect of other types of shocks on structural transformation.

The conceptual framework that structures this debate highlights two main channels that operate in opposite directions. The first channel focuses on the return to productive factors: given a positive productivity shock in agriculture, ceteris paribus the marginal return to factors will increase, increasing demand for labor, capital and land in agriculture. (In general, the literature emphasizes this argument more in terms of returns to labor, but the same principle would apply to the returns to other factors, presuming there is some cross-sectoral factor mobility.) This would generate a positive effect on labor share in agriculture.

The second channel focuses on the non-homotheticity of preferences for non-agricultural goods: a positive productivity shock will increase income and thus will (disproportionately) increase demand for non-agricultural products, assuming the income elasticity of demand is higher for these products. This will generate positive local demand effects for non-agricultural products and stimulate a countervailing shift out of agricutlure (Bustos, Caprettini, and Ponticelli 2016).

Importantly, these early models were based on a closed economy. Subsequent work made the point that the first channel is presumably not relevant in an even partially open economy in which factor prices would be determined by international factor prices (Matsuyama 2009).

More recently, the development literature has identified several other channels through which agricultural shocks may affect structural transformation. First, the alleviation of liquidity and credit constraints: subsistence agricultural households may simply be too poor to switch sectors, especially if switching entails migration or the purchase of a non-divisible productive asset. A positive agricultural shock may thus enable households to amass lump-sum funds to exit; there is some evidence of this phenomenon in the existing literature (Bazzi 2017; Bryan, Chowdhury, and Mobarak 2014; Leight 2020).

Second, effects on input costs: a productivity-enhancing shock to agriculture may reduce the prices of agricultural products that are used as inputs in manufacturing, this stimulating the manufacturing sector. This may be arguably of particular importance in developing countries, where the industrial sector is often concentrated in processing of primary products.

Third, effects on capital supply: if there are frictions in the capital market, then a local positive shock to agriculture may generate more savings and an additional supply of capital that can potentially be directed toward (typically more capital-intensive) non-agricultural production.

Empirical literature: Shocks to the agriculture sector

Rigorously identifying the empirical relationships underlying these theories of structural transformation is challenging given that it requires the conjunction of two elements: a plausibly exogenous shock to the agricultural sector, and the use of a unit of analysis in which agricultural and non-agricultural labor

shares can be measured (i.e., a unit larger than a household or community). This review highlights six relatively recent papers that arguably include both these elements: three analyze positive technological shocks to the agricultural industry (though one does so at the cross-country level); one analyzes an institutional shock; and two analyze climatic shocks. Interpretation of the latter is rendered more complex by the fact that climatic shocks are inherently short-term. The review is not intended to be systematic but rather to highlight particularly important recent contributions in the field.

Table 1 summarizes key characteristics of the papers, including the site, geographic unit, shock of interest, and time period. We also report the estimated semi-elasticity defined as follows: the absolute percentage point decline in agricultural labor share, relative to the percentage increase in agricultural output generated by the shock.

Paper	**Site**	**Geographic unit**	**Shock of interest**	**Time period**	**Estimated semi-elasticity**
Bustos et al. 2016	Brazil	Municipality	New soybean seeds that render soybean cultivation less labor-intensive (labor-augmenting). Secondarily also examining a maize shock that is land-augmenting (allows for 2nd drop annually), but most focus is on soybeans.	1996-2010	.15
Emerick 2018	India	District	Increase in rainfall (associated with positive shock to grain yield)	1998-2013	.15
Colmer 2021	India	District	Increase in temperature (associated with negative shock to agricultural outcomes)	2001-2007	-.5-.6: showing a decline in ag output leads to a decline in ag labor share

McArthur and McCord 2017	Multicountry	Country	Variation in access to fertilizer (driven by interaction of distance to fertilizer production sites, and price)	1960-2001	.23 for low-fertilizer distance country .33 for high-fertilizer distance country
Marden 2017	China	County	Variability in suitability of local area for cultivation of cash crops vis-a-vis grain crops	1985-2005	-.07: as suitability for cash crops increases, the analysis suggests that agricultural employment share is in fact increasing
Hornbeck and Keskin	U.S.	County	Expansion of agricultural aquifer	1920-2002	0. No impact of the agricultural shock on the non-ag economy and no effect on land or labor costs; do not have data on agricultural labor share, but seems plausible to assume this is a zero.

The three papers analyzing technological shocks are Bustos et al., analyzing a technological shock to the soybean sector (and secondarily, to the maize sector) in Brazil at the municipality level; McArthur and McCord, analyzing variation in access to fertilizer at the country-time level linked to distance from fertilizer production sites and shifts in the availability of fertilizer over time; and Hornbeck and Keskin, analyzing variation in access to water at the county level in the historical U.S. linked to the expansion of an agricultural aquifer (Bustos, Caprettini, and Ponticelli 2016; McArthur and McCord 2017; Hornbeck and Keskin 2015). The estimated elasticities vary from weakly positive (.15 in Bustos et al. and around .25 in McArthur and McCord, higher for countries that are more distant from fertilizer production sites) to zero in the case of Hornbeck and Keskin, where they identify no effects of the agricultural shock on the non-agricultural economy or land and labor costs in the historical U.S. Historical data on labor share is not available in their analysis, but the available results suggest a null relationship is likely.

Two papers analyze climatic shocks, both using district level in India over roughly the same period (Emerick 2018; Colmer 2021). The findings are strikingly different: Emerick analyzing rainfall shocks finds a weakly positive elasticity (.15) broadly consistent with the magnitudes described previously, and Colmer analyzing temperature shocks presents evidence of a negative elasticity (i.e., high temperature shocks that lead to a decline in agricultural output also generate a decline in the local agricultural labor share as workers shift out of agriculture), of around -.5. For both papers, it seems reasonable to question whether these estimated elasticities correspond to only very short-term effects. In addition,

the striking difference in the estimated signs for rainfall vis-à-vis temperature shocks is notable, though the two analyses are also framed as analyzing the effect of positive shocks (more rainfall) vis-à-vis negative shocks (higher temperature), and there could be theoretical reasons for asymmetry: i.e., if negative shocks push poor households below some subsistence threshold, they could respond very differently. In addition, Colmer argues that the (negative) relationship between temperature shocks and agricultural output is more robust than the relationship between rainfall shocks and output, and that the latter is not of consistent significance or sign once the specification also includes additional controls for rainfall.

Neither paper engages at length with the question of whether the reallocation across sectors is likely to be short- or long-term, a question that is probably impossible to answer given the empirical strategy employed. Colmer does show additional specifications from a subsample of formal manufacturing firms that suggests that the increase in manufacturing employment in response to an adverse temperature shock is primarily driven by contract workers, a pattern that would be consistent with a temporary response; however, formal firms are only a small fraction of the overall sample (constituting 20% of manufacturing employment).

Finally, a recent working paper by Marden analyzes an institutional shock in China: the reforms linked to decollectivization of agriculture and the institution of the Household Responsibility System in China, a shift associated with a substantial shift from staple to cash crops. The analysis exploits cross-region variation in the suitability of counties for cash crops, an index generated from the FAO GAEZ data. This paper also finds a negative elasticity of estimated value around -.07: as suitability for cash crops increases, the agricultural employment share is also increasing. One important point to note is that while the analysis focuses on the broad comparison of cash and staple crops, both categories include a range of crops that were experiencing a wide range of different shocks in China during this period, and this generates some challenges in interpretation (Marden 2015).

This is clearly a wide dispersion of estimated elasticities: (-.5, -.07, 0, .15, .15, .25). A simple median would suggest the elasticity is around .08. The Bustos et al. estimate can reasonably be considered to be the most convincing given that it is the only paper that analyzes a technological shock that is both plausibly exogenous and evaluated at the local level using detailed data, and that would suggest a slightly higher estimate of around .15. (The Hornbeck and Keskin paper also evaluates a technological shock at the local level, but is limited in the historical data available and thus cannot fully trace out relevant patterns of structural transformation.)

We can also use these constructed elasticities to estimate the relative salience of the direct and indirect positive effects of a positive shock to the agricultural sector. The direct positive effect is simply defined (if agricultural productivity increases, agricultural income increases); the indirect positive effect is generated by the shift of productive factors out of agriculture and into non-agricultural production, a higher-productivity sector. (This exercise is not well-defined for the two cases, Colmer and Marden, in which the elasticity is identified to be negative; and in the case of Hornbeck and Keskin, the indirect effect would clearly be zero.)

The only paper that directly reports this exercise is Emerick, who estimates the indirect positive effect to be a third of the direct effect, or a quarter of the total effect. We can repeat the calculation for Bustos and McArthur and McCord, using estimates of the agricultural labor share at baseline and the estimated gap in labor productivity comparing across sectors; both suggest that the indirect positive effect via

structural transformation of a positive shock to agriculture is around half the total effect.[1] The calculations underlying this exercise are reported here. In general, the larger the estimated gap in productivity across sectors, the more salient the indirect effect will be – as the gains from shifting factors are estimated to be even larger. Nonetheless, this exercise suggests that the gains from any positive shock to agricultural productivity that accrue via structural transformation are, at best, likely equal to the direct positive effect of the shock, and no larger.

Empirical literature: Other types of shocks

In addition to this more well-defined set of papers analyzing productivity shocks to agriculture, there are also four other recent papers analyzing other types of shocks. Dinkelman et al. 2017 analyze variation in local exposure within Malawi to a lump-sum shock corresponding to remittances of earnings from migrant mining workers working in South Africa. Bustos et al. 2020 build on their previous analysis in Brazil to analyze the effect of a local increase in capital supply – savings linked to increased soybean profits. Asher and Novosad (2020) analyze village-level effects of a road-building program; roads could have a range of effects and could potentially increase the returns to factors in agricultural production (due to increased market access, lower input costs, etc.) or non-agricultural production (for similar reasons, and particularly due to the lowest costs of accessing local urban areas) or both. Erten and Leight (2021) analyze county-level effects of reduced tariff uncertainty for (primarily) manufacturing exports in China, a shock that increases the returns to capital and labor in the non-agricultural sector.

Given that these papers are all rather different, it is less feasible to draw out consistent patterns. The analysis in Malawi exploits in the timing and volume of migrant flows to South Africa, leading to varying amounts of remittances that flowed back when migration was abruptly cut-off following a plane crash in 1974 (Dinkelman, Kumchulesi, and Mariotti 2017). The data employed then allows the authors to track the effect of these remittance inflows on structural transformation over the next thirty years; they benchmark their primary results vis-à-vis the effect of a million dollar of remittances, while the average district received $2.25 million. The findings suggest this inflow leads to a 1.5 percentage point shift in labor share in agriculture.

The second paper in Brazil has a more complex design in which the authors identify the increase in bank deposits (derived from agricultural profits) linked to the positive shock to soybean production in each municipality; the analysis then uses data on the full universe of bank branches in Brazil to identify how these shocks propagate through links in the financial system between capital-sending and capital-receiving municipalities, and find that these positive shocks are associated with increased loans to non-agricultural firms in capital-receiving municipalities (Bustos, Garber, and Ponticelli 2020). This paper frames the magnitude relative to a certain percentage increase in soybean profits; using some additional sources, a very rough estimate suggests that a positive shock to soybean profits of a million dollars leads to around a .4 percentage point shift in labor share in agriculture. This is significantly smaller than the

[1] As a general baseline, foundational work suggests that the productivity gap comparing across the non-agricultural and agricultural sectors is around 3.5 for all countries, and 5.6 for the poorest quartile of countries; following adjustments for human capital quality and input intensities, the gap reduces to 2.3 for the full sample, and 3 for the poorest quartile of countries (Gollin, Lagakos, and Waugh 2014). Emerick provides an estimate of an agricultural productivity gap of 2.5 in India; Pavcnik and McCaig estimate a gap of 5 in Vietnam; Dinkelman et al. estimate a gap of 1.12 in Malawi; and Erten and Leight estimate a gap of around 6 unadjusted and 4 adjusted for human capital stocks in China.

effect in Malawi, perhaps unsurprising given that the gap in per capita income comparing across Malawi and Brazil in the period of interest is on the order of $166 versus $5000.[2] Accordingly, Malawian districts may be extremely capital constrained, and experience a much larger shift in non-agricultural investment given the same infusion of cash.

The two other papers identified in this literature analyze local institutional shocks that may shift returns to both agricultural and non-agricultural production (Erten and Leight 2021; Asher and Novosad 2020). Asher and Novosad analyze the construction of roads in rural India, a shock that could potentially increase or decrease labor demand in both agricultural and non-agricultural sectors. They find that road construction leads to a 9 percentage point decline in the local share of workers in agriculture, who have shifted to wage labor in nearby towns; but no other effects on consumption or assets, and only minimal effects on village non-agricultural production. This paper poses something of a puzzle vis-à-vis the other papers in this literature, in two respects. First, while the findings suggest that transportation costs are a barrier to the reallocation of labor, there is no evidence that this reallocation of labor is associated with higher productivity, wages, or consumption, and thus it is unclear what motivates workers to switch (though there could be non-pecuniary returns to switching into wage labor, or anticipated future returns). Second, there is no evidence that this switching is driven by any local agricultural windfall, as the effects on agricultural production are zero.

Erten and Leight analyze the effect of accession to the World Trade Organization in China on local structural transformation. Accession is associated with reduced barriers to exports and thus increased returns to factors in local non-agricultural production, and as expected, this leads to an outflow of productive factors from the agricultural sector and the expansion of the local non-agricultural sector. Here, the implied magnitudes of the shocks are much larger since China is generally growing very rapidly in this period. We find the WTO accession shock is associated with about 10% of GDP growth for the average county over the post-WTO period (2002-2013); GDP per capita increases by about $6000 over this period, implying a potential gain of around $150 million for the average county of population 250,000, and this gain is associated with a decline of about 7 percentage points in the agricultural labor share. This paper was the first to use trade-linked shocks to the returns to factors in non-agricultural production to analyze effects on structural transformation.

Manufacturing, or manufacturing and services?

One recent contribution to the structural transformation literature presented evidence on cross-country sectoral gaps and highlighted two other points relevant to this discussion (Herrendorf, Rogerson, and Valentinyi 2022). First, the benefits in terms of productivity for developing country economies may also be realized by shifting productive factors out of agriculture and into services (i.e., the benefits are not limited to a shift into manufacturing). The paper frames the analysis as focusing on manufacturing productivity, agricultural productivity, and aggregate productivity, but the omitted sector is services, and the key implication is that there may also be gains by transitions out of agriculture into non-manufacturing (by definition, this is services). Second, this new evidence questions the conclusion (previously suggested by Dani Rodrik) that manufacturing productivity is converging across countries over time.

[2] Dinkelman reports the estimate of $166 for Malawi in 1998, and World Bank data suggests that in the same year, Brazil's per capita income in PPP terms was around $5000.

Overall, this paper is broadly consistent with the existing literature that generally highlights the distinction between agricultural and non-agricultural production, rather than specifically manufacturing. While certainly the popular narrative around development highlights manufacturing-driven growth in contexts such as China and Vietnam, in fact the productivity gap as described by Gollin et al. is framed as agricultural versus non-agricultural sectors, and the empirical literature that I've described here generally analyzes shifts out of agriculture into both manufacturing and services. For example, even in Erten and Leight where the shock is explicitly to the manufacturing sector, the realized effects are of equal magnitude in agriculture and services. Thus while the Herrendorf et al. paper may be something of a departure to popular narratives of development that emphasize manufacturing, it is not a significant departure from the underlying academic literature.

On the second point around convergence in manufacturing productivity across countries, from the perspective of developing economies, the key point is whether non-agricultural production is characterized by higher value-added per worker vis-a-vis agricultural production in the same country. All the evidence suggests this is true; and as long as it is true, our interest in structural transformation from the point of view of increasing aggregate productivity in developing countries would remain, even if manufacturing productivity in developing countries is not converging upwards over time.

Summing up

Overall, a few key points can be highlighted from the recent literature on structural transformation.

1) There is relatively little consensus around the elasticity of the local labor share in agricultural production with respect to local agricultural shocks, and little consensus even on what might define a shock (i.e., differentiating between technological and climatic shocks); however, the weight of the evidence seems to suggest a low, positive elasticity.

2) Given this elasticity, it seems reasonable to conclude that the majority of the welfare benefits from any positive shock to the agricultural sector will be constituted by the direct benefits of higher output in the agricultural sector, rather than generating an outflow of factors to the more productive non-agricultural sector; estimates suggest that at most 50%, and possibly less, of the total benefit is constituted by the indirect benefit. Ceteris paribus, a higher productivity gap across sectors will signify a higher indirect benefit, as the local economy will then benefit even more from factors entering a much more productive non-agricultural sector.

   We could hypothetically imagine a scenario in which the indirect effects of a positive agricultural productivity shock are at least four times larger than the direct effect (i.e., the direct effect is 20% or less than the total effect). Assuming the productivity gap comparing across the agricultural and non-agricultural sector is three-fold as suggested by Gollin et al. for the poorest 25% of countries, the semi-elasticity would have to be 1 to generate a this scenario: i.e., the agricultural employment share would have to decline by 1 percentage point for each 1% increase in agricultural output.[3] This parameter value for the semi-elasticity is around seven

[3] More specifically, this calculation corresponds to the following scenario: a baseline agricultural labor share of 50%, baseline value-added per capita in agriculture of $1000, and baseline value-added in non-agriculture of $3000, corresponding to a productivity gap of three.

times the most plausible estimate generated in the literature, and around four times the highest estimates reported in the literature.

3) The literature on other types of shocks (lump sum shocks, or shocks to the non-agricultural sector) is generally not large enough to draw out any systematic empirical patterns. The effect of a lump-sum "cash drop" linked to agricultural profits was much larger in Malawi than in Brazil, possibly reflecting the much lower level of per capita income, and this could suggest that the effects of cash infusions on structural transformation are larger in very poor economies. However, the estimated productivity gap in Malawi is also not particularly high, presumably because household-level non-agricultural enterprises are not particularly productive. Thus while this pattern of structural transformation may have major benefits for household welfare, the aggregate benefits are less clear.

Bibliography

Asher, Sam, and Paul Novosad. 2020. “Rural Roads and Local Economic Development.” *American Economic Review* 110 (3): 797–823. http://www.aeaweb.org/aer/.

Bazzi, Samuel. 2017. “Wealth Heterogeneity and the Income Elasticity of Migration.” *American Economic Journal: Applied Economics* 9 (2): 219–55.

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. “Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh.” *Econometrica* 82 (5): 1671–1748.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli. 2016. “Agricultural Productivity and Structural Transformation: Evidence from Brazil.” *American Economic Review* 106 (6): 1320–65. https://doi.org/10.1257/aer.20131061.

Bustos, Paula, Gabriel Garber, and Jacopo Ponticelli. 2020. “Capital Accumulation and Structural Transformation.” *The Quarterly Journal of Economics* 135 (2): 1037–94.

Colmer, Jonathan. 2021. “Temperature, Labor Reallocation, and Industrial Production: Evidence from India.” *American Economic Journal: Applied Economics* 13 (4): 101–24. http://www.aeaweb.org/aej-applied/.

Dinkelman, T, G Kumchulesi, and M Mariotti. 2017. “Labor Migration and the Structure of Rural Labor Markets.” *Working Paper*. https://www.aeaweb.org/conference/2016/retrieve.php?pdfid=10233&tk=3QYHt779.

Emerick, Kyle. 2018. “Agricultural Productivity and the Sectoral Reallocation of Labor in Rural India.” *Journal of Development Economics* 135 (September 2015): 488–503. https://doi.org/10.1016/j.jdeveco.2018.08.013.

Erten, Bilge, and Jessica Leight. 2021. “Exporting Out of Agriculture: The Impact of WTO Accession on Structural Transformation in China.” *Review of Economics and Statistics* 103 (2): 364–80. http://www.mitpressjournals.org/loi/rest.

Gollin, Douglas, David Lagakos, and Michael E Waugh. 2014. “The Agricultural Productivity Gap.” *The Quarterly Journal of Economics* 129 (2): 939–93.

Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi. 2022. “New Evidence on Sectoral Labor Productivity: Implications for Industrialization and Development.”

Hornbeck, Richard, and Pinar Keskin. 2015. “Does Agriculture Generate Local Economic Spillovers? Short-Run and Long-Run Evidence from the Ogallala Aquifer.” *American Economic Journal: Economic Policy* 7 (2): 192–213. https://doi.org/10.1257/pol.20130077.

Leight, J. 2020. “The Impact of Positive Agricultural Income Shocks on Rural Chinese Households.” *World Bank Economic Review* 34 (1). https://doi.org/10.1093/wber/lhy008.

Marden, Samuel. 2015. “The Agricultural Roots of Industrial Development: ‘forward Linkages’ in Reform Era China,” 1–69. http://www.samuelmarden.com/research.html.

Matsuyama, Kiminori. 2009. “Structural Change in an Interdependent World: A Global View of

Manufacturing Decline." *Journal of the European Economic Association* 7 (2–3): 478–86.

McArthur, John W., and Gordon C. McCord. 2017. "Fertilizing Growth: Agricultural Inputs and Their Effects in Economic Development." *Journal of Development Economics* 127 (July 2015): 133–52. https://doi.org/10.1016/j.jdeveco.2017.02.007.