

DEMAND-SIDE CONSTRAINTS IN DEVELOPMENT:
THE ROLE OF MARKET SIZE, TRADE, AND (IN)EQUALITY

Pinelopi K. Goldberg
Tristan Reed

WORKING PAPER 27286

NBER WORKING PAPER SERIES

DEMAND-SIDE CONSTRAINTS IN DEVELOPMENT:
THE ROLE OF MARKET SIZE, TRADE, AND (IN)EQUALITY

Pinelopi K. Goldberg
Tristan Reed

Working Paper 27286
<http://www.nber.org/papers/w27286>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2020, Revised June 2023

This paper provided the basis of Goldberg’s Presidential Address that was delivered in several regional meetings of the Econometric Society during 2021. An earlier version of this work has appeared as World Bank Working Paper under the title “Income Distribution, International Integration and Sustained Poverty Reduction.” We thank Andrei Shleifer, Michael Keane, Christoph Lakner, Michele Ruta, colleagues at Yale and the World Bank and many participants at the Econometric Society’s regional meetings for useful comments on the earlier version. The current version has benefited substantially from the detailed feedback of the editor and three anonymous referees. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the World Bank and its affiliated organizations, nor those of the Executive Directors of the World Bank or the governments they represent, nor of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Pinelopi K. Goldberg and Tristan Reed. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Demand-Side Constraints in Development: The Role of Market Size, Trade, and (In)Equality
Pinelopi K. Goldberg and Tristan Reed
NBER Working Paper No. 27286
May 2020, Revised June 2023
JEL No. F10,F13,F15,L10,L16,O10,O14

ABSTRACT

What is the pathway to development in a world marked by rising economic nationalism and less international integration? This paper answers this question within a framework that emphasizes the role of demand-side constraints on national development, which is identified with sustained poverty reduction. In this framework, development is linked to the adoption of an increasing returns to scale technology by imperfectly competitive firms that need to pay the fixed setup cost of switching to that technology. Sustained poverty reduction is measured as a continuous decline in the share of the population living below \$1.90/day purchasing power parity in 2011 US dollars over a five year period. This outcome is affected in a statistically significant and economically meaningful way by domestic market size, which is measured as function of the income distribution, and international market size, which is measured as a function of legally-binding provisions to international trade agreements, including the General Agreement on Tariffs and Trade, the World Trade Organization and 279 preferential trade agreements. Counterfactual estimates suggest that, in the absence of international integration, the average resident of a low or lower-middle income country does not live in a market large enough to experience sustained poverty reduction. Domestic redistribution targeted towards generating a larger middle class can partially compensate for the lack of an international market.

Pinelopi K. Goldberg
Yale University
Department of Economics
37 Hillhouse Ave.
P.O. Box 208264
New Haven, CT 06520-8264
and NBER
penny.goldberg@yale.edu

Tristan Reed
The World Bank
Development Research Group
1818 H Street NW
Washington, DC 20433
treed@worldbank.org

1 Introduction

Many of the success stories of economic development during the last century, for instance in East Asia (Stiglitz, 1996), coincided with growth in exports and trade surplus. As advances in technology threaten the comparative advantage offered by cheap labor, and as interest in protectionism rises in advanced economies, it is not clear that this same export-led model will be relevant in the future. Lund et al. (2019) find that the share of trade based on labor-cost arbitrage (defined as exports from countries whose GDP per capita is one-fifth or less than that of the importing country) has been declining in some value chains, especially in labor-intensive manufacturing where it dropped from 55 percent in 2005 to 43 percent in 2017. Such observations have led policy makers to ask: What is the pathway to development in a world with less international integration?

The literature has suggested many alternative theories of development. Rodrik, McMillan, and Sepulveda (2013) provide a useful taxonomy, in which they distinguish between two main schools of thought: one based on theories of “structural transformation” and one based on “fundamentals”. Theories of structural transformation are premised on the idea that in low-income countries, there is a “dual economy” with two sectors, one with high, and one with low productivity, and that development involves the transition from the low- to the high-productivity sector. In contrast, theories based on “fundamentals,” (i.e., human capital, infrastructure, institutions) typically treat the economy as a single sector and emphasize the importance of long-run investments in fundamentals as prerequisite for development.

The present work classifies the ideas highlighted by Rodrik in an alternative taxonomy that distinguishes between theories emphasizing demand-side and theories emphasizing supply-side constraints in development. We consider the two approaches complementary; the evidence certainly provides support for the main ideas of both. However, we argue that for many countries, especially those with small market size, it is the demand-side constraints that are binding.

The importance of supply side constraints to development has been explored in a large literature. Banerjee and Duflo (2005) provide an excellent overview in their Chapter for the 2005 Handbook of Economic Growth. Chief among the supply-side constraints they discuss are capital market imperfections. The recognition that such constraints are important has led to many important policy efforts to alleviate them (including our own work on human capital - see, for example, Angrist, Djankov, Goldberg, and Patrinos (2021), Agarwal and Reed (2022)). But the crucial premise underlying all these efforts is that the demand side of the economy will support supply-side efforts, and hence one does not need to worry about it; if policy makers invest in the “right” long-term policies on the supply-side, the rest will take care of itself. If, for instance, countries invest in

human capital, then there will be jobs for the more skilled labor force such investments will create. Accordingly, the policy focus in this strand of the literature is on the supply side.

In contrast, several theories of structural transformation emphasize the demand side of the market. The basic structural transformation paradigm involves the transition from agriculture to manufacturing and then to services, whereas the development of the service sector is accompanied by urbanization. In this paradigm, the demand side needs to support the growth of the expanding sectors. Openness and export-led industrialization play an important role in this process and ensure that market-size is not a constraint. Large countries (such as China or India) can rely on their domestic markets. Small countries can tap into the international market to increase their effective market size. Exports to the lucrative markets of more advanced economies allow them to exploit economies of scale and obtain revenues that they then can use towards the all-important investments in fundamentals.

Of course, the reality is much more complex than this stylized description of theoretical paradigms as various papers and reviews of structural transformation theories have made clear (see for example the informative discussions in Gollin (2010), Gollin (2014), Gollin, Parente, and Rogerson (2004), and Herrendorf, Rogerson, and Valentinyi (2014)). However, to a first-order approximation, the experience in East Asia (China, South Korea, Vietnam, Thailand) roughly conforms to the paradigm described above.

What is different today? Increasing automation, combined with a strong backlash against globalization and the rise of economic nationalism are making this export-led model of industrialization unlikely. As the old paradigm of structural transformation is coming under strain, demand-side constraints in development are becoming first-order. This leads to the question posed in the opening paragraph: What is the vision for such countries?

One alternative often suggested in policy circles is the growth of services. However, services are a highly heterogeneous category - not all types of services are associated with the high productivity gains that go hand in hand with the growth of manufacturing. Gollin, Jedwab, and Vollrath (2016) for instance, make a distinction between “production” and “consumption” services and show that only the first are associated with productivity gains. In contrast, the growth of consumption services led to “urbanization without industrialization” —a phenomenon particularly salient in several African cities. Closely related is the “premature de-industrialization” of Africa (Rodrik, 2016). In a related vein of work, Fan, Peters, and Zilibotti (2021) ask, this time focusing on India, whether growth in India has been service-led or service-biased and find the answer to this question to be nuanced. In sum, there is little evidence to date that the service sector as a whole can served

as the engine of growth for low- and lower-middle-income countries.

In the quest for a new “vision” for development, the following observation may be useful. India grew fast in the last three decades, but its growth was not driven by manufacturing, and certainly not by exports. This leads to the hypothesis that a large domestic market size, like the one India enjoys thanks to its large population, may make it easier to develop. The reason is that a large domestic market allows countries to exploit economies of scale. But Nigeria, a country both large and rich in natural resources, provides a counterexample. Nigeria has experienced periods of fast growth, but no sustained development and no poverty reduction. There are of course many differences between India and Nigeria, but one particular hypothesis that comes to mind in the comparison of the two countries is that development also requires a certain degree of “equality,” so that a positive shock (e.g., through an increase in agricultural productivity or oil exports) can trickle down, generate demand, and jump-start the process of structural transformation. When the wealth of a nation is held in the hands of a few, this process never takes off. Note that in this story, a certain degree of equality is prerequisite for development.

While the above hypothesis may have some intuitive appeal, it is far from the point where it can serve as the basis of scientific inquiry. Several questions come to mind: How do we operationalize this idea? What does “large market” mean in this context, i.e., how large does a market need to be for domestic economies of scale to be sufficient? What is the meaning of “equality” here? How do we measure “development”? What is a small country supposed to do if it cannot rely on its domestic market and if trade is not an option? And many more...

To answer these questions and bring them to the data, one needs a conceptual framework. To this end, we consider a stylized model in the spirit of Dave Fultz’s “dish-pan” model of global weather patterns, eloquently described in Paul Krugman’s “The Fall and Rise of Development Economics” (Krugman, 1994). As Krugman notes, we know that the model is wrong— how could a model ever do justice to the complexity of the development process? Nevertheless, the stylized framework is useful in helping us to structure our thinking, cut away the extraneous in order to gain insight into the mechanisms through which a larger market size, domestic or international, can spur development, and answer the questions posed above.

Specifically, we develop a model that emphasizes the role of demand-side constraints on national development, which we identify with sustained poverty reduction. We measure sustained poverty reduction as a continuous decline in the share of the population living on less than \$1.90 PPP per day in 2011 US dollars, over a 5 year period. In this framework, development is linked to the adoption of an increasing returns to scale technology by imperfectly competitive firms, who need

to pay the fixed setup cost of switching to that technology. Poverty is reduced as adoption of the new technology sets off a structural transformation process that increases wages. The necessary demand to overcome the fixed cost of technology adoption may come from either the domestic or the international market. Importantly, economies of scale can be achieved even in sectors serving primarily domestic demand (e.g., services); in such a case, the role of exports from the tradable sector is that they provide income that translates into additional demand for all sectors, including those that are non-tradable. Increased demand stemming from broad household ownership of firm profits is the key channel through which an equitable distribution of wealth and income improves labor productivity in the general equilibrium model we use to motivate our empirical analysis (Murphy, Shleifer, and Vishny, 1989).

The size of the international market is measured on the basis of a new database of the legally-binding provisions of international trade agreements, which include the General Agreement on Tariffs and Trade (GATT), the various agreements of the World Trade Organization (WTO), and 279 preferential trade agreements (PTAs) whose provisions are recorded by Hofmann, Osnago, and Ruta (2017). These provisions primarily establish rights related to goods and services trade, but also relate to flows of capital, ideas, and labor, and together make up the legal architecture of the international economy. The size of a given country's integrated international market is calculated by summing the population and income (or, alternatively, size of middle class) of all other countries, where those countries are weighted by the number of economic integration provisions a country has signed with them. Though the relative per capita income of the integrated market declines rapidly with a country's national income in our sample, we find that its effect on sustained poverty reduction is positive, quantitatively large and statistically significant, suggesting international integration has provided a pathway for poor countries to eliminate poverty in the past. These results are highly relevant in light of the existing evidence that poverty declines more slowly among countries that are initially poor (Ravallion, 2012).

Setting the size of the integrated international market to zero in a counterfactual scenario allows us to isolate the effect of domestic market size alone on sustained poverty reduction, which quantifies the hope for development in a less integrated economy. The size of the domestic market is measured as a function of the income distribution, as summarized by the share of the population in the global middle class (defined based on Kharas (2017) to include those living on \$11-110 PPP per day in 2011 US dollars). The use of an absolute, rather than relative, definition for the middle class reflects the assumption that the increasing returns technology is the same across countries. Like the size of the international market, the middle class share has a positive and significant effect on sustained

poverty reduction.¹

Our empirical framework is inspired by the industrial organization literature, specifically Bresnahan and Reiss (1991), who develop a method to estimate entry thresholds based on the profit functions of firms facing increasing returns and entering imperfectly competitive markets. Their approach is especially useful in our setting because it does not require data on market prices to estimate variable profit and fixed cost parameters, which are required to calculate the break-even point. Using this approach, we estimate that the threshold market size for sustained poverty reduction is 325 million people, if the purchasing power of these people is below that of the global middle class.

In a scenario in which the size of the integrated international market is set to zero, as of 2011-15, the average resident of a low and lower-middle income country does not live in a market large enough to experience sustained poverty reduction. The primary reason for this is that the middle class in these countries is not yet large enough. In our preferred specification, for the average country in our sample, increasing the share of the population in the middle class by 10 percent is equivalent to increasing population by 50 million people. For countries with small population, equality therefore is disproportionately important. This suggests that, if international integration is indeed waning, to eliminate poverty policy-makers in poor countries must focus on equalizing the distribution of income, for instance through taxation or (as suggested by the model underlying this paper’s analysis) redistribution of equity shares to the poor.

Methodologically, our work is related to a specific approach in the economic growth literature, in which researchers identify a set of countries that perform exceptionally well over some time period, and then compare them with the rest of the world. In the report of the Commission on Growth and Development, Spence et al. (2008) identify 13 economies that have sustained cumulative GDP growth of more than 7 percent annually for 25 years or more since 1950.² Werker (2012) studies all countries that achieved double-digit growth—above 10 percent annually—for 8 or more years,

¹Allowing for a large gap between the lower bound of the middle class (\$11 PPP per day) and the upper bound of poverty (\$1.90 PPP per day) ensures the relationship between the share of the population in these two states is not mechanical. Alternative definitions of the middle class have been studied by others. For instance, Birdsall, Graham, and Pettinato (2000) and Easterly (2001) define the middle class in each country relative to the national income distribution (i.e., respectively, 0.75-1.25 of median income; and the 20th and 80th percentile in consumption). Banerjee and Duflo (2008) and Ravallion (2009) examine alternative definitions which are fixed across countries, but may be considered suitable for different income levels (i.e., respectively, \$2-4/day and \$6-10/day; or a “developing world’s middle class” with income above the median poverty line of developing countries and a “Western middle class,” above the poverty line of the United States). The use of an absolute threshold anchored to advanced country living standards, as in this paper, is more common in the private sector, for instance among retailers considering whether to enter a market. For the argument that this approach is profit-maximizing, see the critical review of “bottom-of-the pyramid” retail strategies by Simanis (2012).

²These are Botswana; Brazil; China; Hong Kong SAR, China; Indonesia; Japan; the Republic of Korea; Malaysia; Malta; Oman; Singapore; Taiwan, China; and Thailand.

finding that almost two thirds of such periods are either recoveries from war or resource booms, typically those driven by oil. Hausmann, Pritchett, and Rodrik (2005) identify periods of growth accelerations, by identifying all periods in which the change in growth rate is greater than or equal to 2 percentage points per annum, and then coding the successive 7 year period to equal one (and zero otherwise) if growth over that time was more than 3.5 percent per annum, and if income at the end was higher than the maximum of income during the period. They find that such accelerations are highly unpredictable. In all of these studies, the years over which growth is observed are allowed to vary, and the length of time studied is longer than five years. In contrast, our approach holds periods fixed in time, each comprising a disjoint five-year window (e.g., 1981-85, 86-90, etc.). This approach constrains us from selecting windows of time that paint a disproportionately positive or negative picture of performance in a specific country. It also means our predictions are relevant for the relatively shorter time horizon over which governments make decisions.

The paper proceeds as follows. In Section 2, we introduce our conceptual framework that guides the empirical strategy laid out in Section 3. In Section 4, we describe the variables we construct to bring the model to the data, namely sustained poverty reduction, the middle class share of the population, and the relative size of the integrated international market. Section 5 presents the results. Section 6 discusses our counterfactual estimates of market size in an economy without international integration. Section 7 concludes and offers some thoughts on policy implications.

1.1 Related Literature

In addition to the works just described, our analysis contributes to several distinct literatures. First, our focus on sustained poverty reduction is relevant to a literature on poverty dynamics, which have been studied in individual countries (Ferreira, Leite, and Ravallion, 2010) and among households (Carter and Barrett, 2006; Baulch and Hoddinott, 2000). A key message of this literature is that households frequently move in and out of poverty, and it is much rarer to escape permanently than to escape for a few years (Shepherd and Diwakar, 2019). Looking across countries, more than half of the time countries have sustained poverty reduction at the aggregate level. The results also highlight the limited effect of the business cycle in advanced economies on poverty reduction in developing economies, at least during the 2006-10 and 2011-15 windows, which included the advanced economies' financial crisis and deceleration, and yet were among the best years for sustained poverty reduction.

Second, our paper contributes to a voluminous literature on inequality, poverty and growth. We find that a certain degree of equality and poverty reduction go hand in hand at low income levels,

a result that is broadly consistent with Barro (2000) and Barro (2008), Keane and Prasad (2002), and Ostry, Berg, and Tsangarides (2014). Our work in this regard is most closely related to Desai and Kharas (2017) who emphasize the importance of the middle class in poverty reduction. While these authors use historical data since 1870 to explore the relationship between the middle class and poverty reduction, we focus on a more recent period that is characterized by growing global integration and use counterfactual simulations to quantify the role of the middle class in sustained poverty reduction.

Third, our study contributes to a literature on the effects of trade policy on poverty (Autor, Dorn, and Hanson, 2016; Topalova, 2010; Harrison, 2007; P. K. Goldberg and Pavcnik, 2004; Winters, McCulloch, and McKay, 2004). We introduce to this literature a novel measure of integrated international market size, which in our model predicts sustained poverty reduction. This measure complements and extends the data sets of Sachs and Warner (1995) and Wacziarg and Welch (2008), which identify the specific years at which economies liberalize. By our treaty-based measure of liberalization, in which a country has access to some international market once it signs a trade agreement, many countries appear open in years when these other datasets consider them closed. Even though many countries are closed according to our measure in 1981, almost none are closed today, given almost complete membership in the WTO by UN member states. Despite this, there is large variation in the relative size of the integrated international market in many regions, driven for instance by China's entry into the WTO in 2001, which lowered the average income of WTO member states by approximately 18 percent.

Finally, while models of development with firm-level increasing returns assume supply-side constraints, such as capital market imperfections, to explain why countries remain poor (Banerjee and Duflo, 2005), our framework does not rely on such assumptions. Our demand-side framework implies that the small size of the market itself may explain why countries remain poor. Support for the view that demand-side constraints may be binding comes from the empirical literature on the growth of small and medium-sized enterprises (Woodruff, 2018). While a decade of research on supply-side interventions, for instance microfinance (Banerjee, Karlan, and Zinman, 2015) and business training (McKenzie and Woodruff, 2014), has found mostly disappointing effects, several studies find that that expansion into export markets allows firms to increase their scale and adopt new technologies (Lileeva and Trefler, 2010; Bustos, 2011; Aw, Roberts, and Xu, 2011), and that more generally, boosts to demand may be effective in promoting productivity growth (Ferraz, Finan, and Szerman, 2015; Atkin, Khandelwal, and Osman, 2017; Alfaro-Urena, Manelici, and Vasquez, 2020).

2 Conceptual Framework

We define “development” as sustained poverty reduction. While many indicators summarize a country’s progress, poverty reduction is arguably the best indicator that a country is on track to becoming what could be called an advanced economy. Poverty elimination is the first of the World Bank’s Twin Goals and the first of the United Nation’s Sustainable Development Goals. All advanced economies have eliminated extreme poverty. For practical purposes, the World Bank defines extreme poverty elimination as occurring when the headcount of people living on \$1.90 per day falls to less than 3 percent of the population, recognizing that some small pockets of poverty will always remain, even in advanced economies. According to the World Bank (2020), the extreme poverty headcount in the United States is 1.25 percent, in Japan, 0.22 percent, and in Germany, 0 percent.

Our focus on the transition between two dichotomous stages of development, one with extreme poverty and one without, follows in the tradition of W. Arthur Lewis and others. In this framework, the economy has two alternative production technologies, one with constant returns to scale and another with increasing returns to scale.³ Development occurs when firms pay the fixed setup costs of adopting the increasing returns technology, which causes labor productivity to rise. Even if the poor do not work in the firms that adopt the new technology, poverty falls because the common wage paid to all workers rises.

The main implication of this framework is that a *threshold market size* is required to achieve development—if there is not enough demand, a firm operating the increasing returns technology will not break even. Development is given by the threshold crossing model

$$D = \mathbb{1}(\Pi > 0) \tag{1}$$

where Π is profitability in the increasing returns sector.

The idea that international markets allow firms to achieve minimum efficient scale is well established in trade theory (Helpman and Krugman, 1985). In principle however, a large enough domestic market could also allow firms using the increasing return technology to break even. Murphy, Shleifer, and Vishny (1989) provide a model of exactly this phenomenon, with a specific mechanism in which the effects of a positive income shock, from either agricultural productivity or exports, depend on the initial shareholdings of individuals in society. Societies develop faster when shares in the firms are distributed more equitably across the population, raising the marginal propensity to consume

³Banerjee and Duflo (2005) propose a similar model of development in which firms choose to upgrade to a new technology, and emphasize the role played by capital market imperfections in prohibiting the adoption of this technology.

out of the profits generated by the increasing returns sector⁴. This model suggests that a large internal market may provide a path to development, even in the absence of trade. The middle class, which determines the size of this market, is the result of an initial wealth shock *and* an initial relatively equitable distribution of firm ownership.

3 Empirical Strategy

Bringing a highly stylized model to the data is challenging. As noted earlier, we know that the model is “wrong.” Accordingly, we do not attempt to test it against alternative hypotheses. Instead, our empirics are very much in the spirit of Leamer and Levinsohn (1995) who in their Chapter for the Handbook of International Economics encourage empirical researchers to “estimate, not test.” We use an empirical model inspired by the theoretical framework described in the previous section to estimate the threshold market size required to achieve sustained poverty reduction, and answer the questions posed in the Introduction related to the role of domestic and international market size in development. In line with this thinking, the empirical model includes only variables suggested by the theoretical framework.

Before we present the empirical model in detail, a couple of remarks are in order. First, the theoretical model is static in nature, but makes statements about an inherently dynamic process, the process of development. The estimation utilizes cross-country panel data for countries in which extreme poverty still exists at the beginning of the sample period.⁵ The specifics of the empirical model will be explained in the next section, but we note here that identification of the parameters relies heavily on cross-country variation, as most of the key variables in the analysis (population, various measures of market size) move only slowly over time. However, they do move (otherwise there would be no development, i.e., poverty reduction), so that the time dimension of the data is essential for identification. Given the slow movement of most variables and the fact that our focus is on *sustained* poverty reduction, and not on transitory changes, we do not utilize annual data, but employ 5-year windows instead.

Second, we do not measure increasing returns to scale by estimating production functions.⁶ Our

⁴In a different context, Keane and Prasad (2002) provide cross-country evidence from the transition economies in the 1990s showing that domestic redistribution that reduced inequality promoted growth.

⁵Given that the focus is on the forces that set off the process of development, and not on the final outcome, we exclude countries that are already “developed” during our sample period. Developed countries may look today very different from when they started to develop, and one would need to go back in history and analyze historical data in order to understand their paths.

⁶Returns to scale have been empirically elusive (some of the reasons why empirical work may fail to find evidence for them even when they exist go back to the Klette and Griliches (1996) critique of revenue-based production function estimation). However, a recent literature has found indirect evidence for scale economies in several settings based on

empirical framework allows us to assess whether the results are consistent with the presence of increasing returns, and if so, quantify their implications for the threshold market size.

3.1 Empirical Model

Our empirical model is based on the threshold crossing model of Equation 1. The goal is to estimate a profit function allowing for increasing returns, so that we may calculate the relative contribution of domestic and international market size to sustained poverty reduction. Bresnahan and Reiss (1991) propose a method to estimate the profit function of a profit maximizing firm when data on prices and quantities are unavailable. We adopt their approach to modeling the profit function, while letting the dependent variable $D = D_{it}$ be an indicator that sustained poverty reduction is achieved in country i over the five-year period indexed by t .

Suppose profit of the increasing returns sector in country i at time t is given by:

$$\Pi_{it} = S(M_{it}, \lambda) \times V(X_{it}, \beta) - F(\gamma) + \epsilon_{it} \quad (2)$$

where λ , β and γ are parameters of the market size, variable profits, and fixed costs respectively, M_{it} are observables capturing market size, X_{it} are per-capita demand and cost shifters determining variable profits, and the error term ϵ_{it} captures unobservable factors affecting profits.

Variable Profits Equation (2) corresponds to the functional form of Murphy, Shleifer, and Vishny (1989), in which expenditure of the middle class is multiplied by profits generated by an income boost due to agricultural productivity or export growth to determine the level of industrialization. The model that opens their paper postulates a closed economy. If we were to take that version of the model literally, we would include only agricultural productivity as a measure of the income boost. But no economy is completely closed. Even North Korea has some trade with China. We therefore follow an extended version from Section V of their paper, which allows for a specific, very limited form of trade: the “poor economy” exports commodities or basic manufactures, and these exports play the same role as agricultural productivity in generating an income boost. We assume:

$$\begin{aligned} V &= X_{it}\beta \\ &= \beta_1 + \beta_2 \text{ export growth}_{it-1} + \beta_3 \text{ agricultural total factor productivity growth}_{it-1} \end{aligned} \quad (3)$$

the response of production quantities to plausibly exogenous demand shocks (Costinot et al., 2019; Bartelme et al., 2019).

where $t-1$ refers to growth over the preceding five year period. The variables included in X_{it} account for differences in the variable per-capita profits of the increasing returns sector across countries.

Market Size This extension is more realistic, but still presents the problem that all countries have some degree of integration with foreign economies—as evidenced by the signing of trade agreements, which is our measure of international integration. This calls for including some measure of the integrated international market when specifying the market size. Their model gives no guidance as to how to deal with partial international integration, so some judgment calls are necessary.

The function S summarizes the domestic and international market as determined by population, the income distribution, and international integration. We assume a linear function:

$$\begin{aligned}
 S(M_{it}, \lambda) = M_{it}\lambda = & \text{population}_{it} + \lambda_1 \text{ middle class share of population}_{it} & (4) \\
 & + \lambda_2 \text{ relative population of integrated market}_{it} \\
 & + \lambda_3 \text{ relative GDP per capita integrated market}_{it}.
 \end{aligned}$$

The domestic market size is specified as a function of the domestic population and the share of that population that is in the global middle class. We set the coefficient of population in $S(M_{it}, \lambda)$ equal to one because V contains a constant term. This normalization translates units of market demand into units of population, allowing for an easy interpretation of S . So that our quantitative estimate of market size is more easily interpretable in terms of people consuming less than the middle class, before estimation we subtract from population the number of people in the middle class.⁷

The international market size is specified as a function of the relative population of the integrated international market (relative to the domestic population) and an additional measure that captures the scale of profits generated by selling to the integrated market. In constructing this additional measure, we consider two alternatives: the share of the integrated market that is in the global middle class, and the GDP per capita of the integrated market relative to the GDP per capita of the domestic market.

Both measures (share of global middle class and relative GDP per capita) can be justified based on different assumptions about who bears the fixed cost of the increasing returns to scale technology.

⁷This decision does not substantially affect our estimate of λ .

Our preferred measure is the relative GDP per capita for the following reason. In the closed economy version of the model, the upper class of the domestic economy pays the fixed cost of the increasing returns to scale technology. If the good were sold only to the upper class, the technology would just break even, yielding zero profits. The domestic middle class generates the profits. The larger the middle class, the larger the profits, and the stronger the impetus to industrialization. If we assume that in the partially open economies of our sample the fixed cost is still paid by the domestic upper class (and not by foreigners), then any sales to consumers other than those belonging to the upper domestic class generate profits. These consumers include the domestic middle class (as in the closed economy case) and all foreign sales (both the sales to the foreign middle class and the sales to the foreign upper class). This calls for measuring international market size using all foreign consumers who can afford the increasing returns to scale product, so that income per capita of the foreign market (relative to the domestic one) is the appropriate measure of the international market as it captures the purchasing power of the international market.

Alternatively, one could assume that the fixed cost of the increasing returns to scale technology is borne by both domestic and foreign upper class consumers. In this case, the proper measure would include the share of the integrated market that is in the global middle class. This specification presumes integrated capital markets, so that foreign investment can be used to fund the fixed cost. In contrast, the relative GDP per capita specification does not require any capital to flow between countries: the increasing returns to scale technology is funded only by the domestic upper middle class, but the product produced by this technology is sold both to the domestic and to the foreign market. This latter assumption may be a more realistic description of the early stages of industrialization (usually capital does not flow into poor countries until development has taken off).

One additional consideration is that the effect of agricultural productivity on variable profits in V may vary with the degree of openness. While agricultural productivity provides a boost to income in a closed economy, in the “Dutch disease” open economy model of Corden and Neary (1982) and Matsuyama (1992), a Hicks-neutral increase in agricultural productivity can also boost comparative advantage in agriculture and potentially retard industrialization. For this reason, we also estimate an alternative version of the model where the coefficient on agricultural productivity growth can differ if the economy is closed, as indicated by the fact that it has not signed any international trade agreements. A further complication is that in an open economy setting, the effect of agricultural technical change on industrialization depends on the factor bias of technical change. Bustos, Caprettini, and Ponticelli (2016) show that the effect can be positive if technical change is labor-saving. To examine whether factor bias of technological change in agriculture affects

our results, we also experiment with specifications that use labor productivity in agriculture (rather than Hicks-neutral technical change) as a control.

Fixed Costs We model the “fixed cost” term as a constant $F(\gamma) = \gamma$. This postulates a model in which the increasing returns to scale technology is the same across countries, and allows for a simple test for the presence of increasing returns, i.e., whether $\gamma > 0$.

It is possible to also include in the fixed cost term controls for long-run determinants of institutions or technology. For instance, it is well known that tropical countries have had poor long term economic performance due either to geographic disadvantage (Sachs, 2001), or interaction with historical shocks, such as colonialism (Acemoglu, Johnson, and Robinson, 2001). However, controls such as tropical share of land area would function similarly to fixed effects and risk overfitting, which is especially easy to do in a discrete dependent variable model. For this reason, we do not include institutional and geographic factors in our model, but rather, in a robustness analysis, compare the fit of our model to the fit of a model in which sustained poverty reduction is determined only by institutional and geographic factors.

The further assumption that ϵ_{it} is normally distributed with mean zero, combined with the threshold condition in Equation (1), yields a probit model in which the probability of development, conditional on market size, demand and costs, is:

$$\Pr(D_{it} = 1) = \Pr(\Pi_{it} > 0) = \Phi(\bar{\Pi}_{it}) \quad (5)$$

where D_{it} is an indicator variable equal to 1, if there is sustained poverty reduction (our measure of development) in country i over period t , and zero otherwise, $\bar{\Pi}_{it} = \Pi_{it} - \epsilon_{it}$, and $\Phi(\cdot)$ is the normal cumulative distribution function. We estimate this model using maximum likelihood.

Threshold market size Estimated threshold market size is given by

$$\hat{S} = \frac{\hat{\gamma}}{\bar{\mathbf{X}}\hat{\beta}} \quad (6)$$

where the overbar indicates sample averages and the circumflex indicates estimates from the probit model. By setting $S(M_{it}, \hat{\lambda}) = \hat{S}$, it is possible to determine which counterfactual combinations of M_{it} would be sufficient for a country to achieve development.

3.2 Descriptive Statistics

To estimate the empirical model, we construct three new variables and also exploit variables from existing sources. In this subsection, we describe the construction of each variable used in the

empirical model, and how these variables differ on average between periods of sustained poverty reduction and periods without sustained poverty reduction. Further details on construction of the variables and additional summary statistics are in the Online Supplement.

All variables are predetermined in each of our data windows, as t refers to the first year of the five-year period. Our sample of 377 country-five-year periods includes 94 countries, observed in 5-year periods from 1981-2015, where five year periods range between 1981 to 2015: 81-85, 86-90, 91-95, 96-00, 01-05, 06-10, 11-15. This sample excludes advanced economies, or those with a poverty headcount below 3 percent for all years in the data. Advanced economies cannot provide any information about sustained poverty reduction, because poverty has been eliminated (following the World Bank definition) for all years they are observed.

Sustained poverty reduction The first key variable is our outcome, D_{it} , a binary variable indicating whether a country experienced sustained extreme poverty reduction over a five-year period. The data used to construct our outcome variable are from the World Bank (2020), which reports the national extreme poverty headcount, or the percent of population living below \$1.90 PPP in 2011 US dollars, in years where household survey data are available. The indicator variable has the advantage of a clearer link between the empirics and the theoretical model, which makes qualitative statements about industrialization/development, but is silent on magnitudes. It also allows us to identify periods of “*sustained* poverty reduction” as 5-year windows during which the poverty count is consistently reduced abstracting from cases where poverty reduction is transitory due for example to a short-lived commodity price boom.⁸ Figure 1 shows this variable for select countries.

To summarize the main patterns in the poverty reduction data, Figure 2 displays the share of five-year periods which achieved sustained extreme poverty reduction, over time across the world and five continents. Notably, the share is above 50% for most years in most continents. This reflects the tremendous progress that has been made against extreme poverty in the last four decades. For instance, sustained poverty reduction occurs in more than 70 percent of countries in 06-10 and 11-15 in Africa and Asia, and in 06-10 in Europe. Interestingly, the period of 06-10, which spans much of the recent crisis originating in the United States, does not seem to have been a particularly bad year for poverty reduction globally. This highlights a separation of cycles between advanced and developing economies. Two eras of poor performance stand out, 1981 to 1995 in Africa, and 1985

⁸On the other hand, the use of the indicator variable means that we lose variation in poverty reduction rates. We examine the relationship between these poverty reduction rates and our key explanatory variables in Table A2 in the Appendix and show that this relationship is in line with the mechanisms postulated by the theoretical framework. However, these results are not useful in inferring threshold market sizes.

to 2005 in South America. Sustained poverty reduction in sample countries in North America has also been limited, except in 11-15.

There is a strong presumption in economics that growth and poverty reduction go hand in hand. To examine whether this relationship bears out in our data, Table 1 compares instances of sustained poverty reduction to instances of sustained growth in real GDP per capita, the latter of which is considered to occur when real GDP per capita does not contract at all during a period. To prevent higher frequency observations of GDP per capita from affecting our results, we look only at GDP per capita in the years in which the poverty headcount is also measured, and interpolate GDP per capita between years, as we have done for the headcount. What is apparent from Table 1 is that sustained GDP per capita growth does seem to be positively correlated with sustained poverty reduction, with 76 percent of instances of poverty reduction also having sustained GDP per capita growth, and only 24 percent not having it. More interestingly however, despite this positive relationship, there are 52 percent of the periods without sustained poverty reduction experience sustained real per capita GDP growth. Poverty reduction and GDP growth are positively related, but growth by itself does not always imply poverty reduction.

Domestic market size The second variable we construct is a measure of domestic market size, which depends on the income distribution. Consistent with the underlying conceptual framework, we identify “domestic market size” with the share of the population in the global middle class. We define the global middle class following Kharas (2010) and Kharas (2017), who proposes bounds at \$11-110/day PPP in 2011 US dollars of consumption, on the basis that the lower bound is the average of the national poverty lines in Portugal and Italy, and the upper bound is twice the median income in Luxembourg. That is, to be in the global middle class, one cannot be poor in the poorest rich countries, and cannot be rich in the richest country. We apply these bounds and estimate the share of the middle class as the headcount ratio for the upper bound minus the headcount ratio for the lower bound.

This middle class share variable combines measures of both equality and average income. To measure equality, we use the average Gini coefficient for each country reported (Milanovic, 2013). For average income, we use data on real GDP per capita from the Penn World Tables 9.1 (Feenstra, Inklaar, and Timmer, 2015). A difference between our approach to measure the global middle class and that of Kharas (2010) is that he uses average household consumption for average income, whereas we use GDP per capita. GDP includes expenditure in the investment sector (i.e. construction, machinery and equipment), thus increasing average income and making our estimates of the

middle class larger relative to Kharas (2010), who focuses on the consumer market. As a result, we have not made an assumption about whether increasing returns are differentially available in the investment or consumption sector. More generally, GDP is preferred to income based on household surveys given the risk of top-coding, and the fact that we are interested in the middle and top of the distribution, especially in low and lower-middle income countries (Deaton, 2005; Ravallion, 2003). The use of national accounts in place of average income is most controversial when studying the lower tail of the income distribution (Pinkovski and Sala-i-Martin, 2014).

International market size The third variable is a measure of relative international market size, according to legally-binding provisions to international trade agreements. We define each country’s global market by summing up the populations and incomes of all other countries, with weights corresponding to the number of legally enforceable provisions of multilateral agreements between the country and all others. Table 2 lists the 32 legal provisions in our data set, and shows that they fall into two broad groups, which together make up the legal architecture of the international economy.⁹

This measure has three main advantages for our analysis. First, it allows us to measure directly the effect of international integration treaties, participation in which is a policy choice for the government. Other analyses of trade liberalization (Sachs and Warner, 1995; Wacziarg and Welch, 2008; Easterly, 2019) have typically focused on a mixture of policy decisions (e.g., liberalizing state monopolies in exporting sectors), and trade outcomes (e.g., abnormally low shares of trade to GDP). By focusing specifically on the policy decision to integrate economically through trade agreements, we ensure our counterfactuals are tied to policies actually within government’s control. Second, because we calculate market size in terms of GDP and population, these measures allow us to estimate directly the relative value of integrating with a richer versus a more populous market. Finally, our measures allow us to exploit variation in market size stemming from the entry of *other* countries into a trade agreement. A good example of this variation is what may be called the China shock to the WTO, shown in Figure 3. When China entered in 2001, GDP per-capita of WTO member states fell from above \$11,000 to below \$9,000, while population increased by more than

⁹We focus on what are called “core” provisions, those related directly to trade (Baldwin, 2008). Non-core provisions cover a wide variety of topics, for instance related to the enforcement of human rights, labor or environmental standards, as well anti-money laundering, consumer protection, and statistics cooperation. Our exclusion of non-core provisions has practical implications for measurement. Since we weight countries by the number of provisions, we do not want to overweight regional trade agreements, which include many more non-core provisions relative to the international agreements. We also experiment with an even narrower definition of international integration, according to which only provisions covering trade in goods and services are included (so we exclude provisions covering labor and capital flows). However, our preferred measure is the one covering all flows related to trade (directly or indirectly), as typically, such provisions are negotiated jointly in trade agreements.

1 billion people. Below we describe in detail how we construct our relative size measures, and, provide an example of the variation we exploit in our estimation by tracing out the China shock to the WTO through the relative population and income of countries' international markets.

Figure 4 displays these variables, averaged in each year for the world and within continents, where observations are weighted by the population of each country. The figure allows us to see how the relative size of the international integrated market in each region changes over time. Each line is a population weighted average of the relative market size, in that year. A number of observations stand out. First, Africa was an early integrator, with many of its largest economies joining GATT early on, for instance South Africa (June 13, 1948), Nigeria (November 18, 1960) and Kenya (February 5, 1964). For a period between the 1970s and 1990s it had the largest relative market size in terms of population. Second, Africa, Asia, and to a lesser extent South America have both integrated with significantly richer countries. Until the 2000s, Africa experienced rapid growth in the market size to which it was linked, as more rich countries joined trade agreements such as GATT. Then, in 2001, when China enters the WTO, relative income falls. Countries in Africa, and other continents, no longer had simply open markets with rich buyers, but a rival in their income bracket. This change in countries' income-based international market size may have had important implications for development—a point also made in a recent paper by Atkin, Costinot, and Fukui (2021), who argue that China's entry in world markets pushed many countries, especially in Africa, towards the bottom of the development ladder. Chiquiar and Tobal (2019) similarly show that Mexico reallocated production away from skilled production and towards unskilled production as China's comparative advantage in skilled production emerged post 2001. Third, countries' own per-capita GDP and population growth affect the relative size of the market. This can be seen clearly in Asia, where the relative market size in terms of GDP per capita declines over the 2000s, as China gets richer. It is also possible to see how the relative population of Africa's international market declines in the 2000s, as Africa's population grows faster than the rest of the world's.

Variable profit We use existing data sets to measure the components of variable profit. To calculate export growth, the first determinant of variable profit, we use the Penn World Tables 9.1 export series (Feenstra, Inklaar, and Timmer, 2015). For total factor productivity growth in agriculture, the second determinant of variable profit, we use the international agricultural productivity series of the United States Department of Agriculture's Economic Research Service (Fuglie, Wang, Ball, et al., 2012).

Differences between periods of sustained poverty reduction and periods without sustained poverty reduction Table 3 presents descriptive statistics for all variables, as well the difference in means between samples and its standard error. Beginning with population, it is clear that periods of sustained poverty reduction occur in larger countries, with 60 million more people on average. The middle class is also larger at the onset of periods of sustained poverty reduction. These results provide some initial support for our hypothesis that domestic market size matters. Similarly, the integrated international market size of countries that experienced sustained poverty reduction tends to be larger than the integrated international markets size of countries that did not, irrespective of whether it is measured by relative population, relative income per capita, or the share of the integrated market in middle class. These results are in line with our framework.

Turning to the income boosts, we find that earlier export growth is significantly higher in instances of sustained poverty reduction, 10 percent on an annualized basis over the last 5 years, compared to 6 percent in periods without. In these simple differences in means, effects of international markets on poverty are loaded both on exports and on international market size. In our structural estimation of the profit function, we will study their effects when they are both included in the same model. Turning to agriculture, there is no significant difference in agricultural productivity growth, which is 1 percent annually in both samples.

4 Results

We now turn to our estimates of the threshold model, which are reported in Table 4. Each column of the table reports coefficient estimates of each parameter, as well as our estimate of \hat{S} , the threshold market size required for the increasing returns sector to break even. The various columns show results for alternative specifications of the international market size and variable profit components of the model, as well as for alternative definitions of sustained poverty reduction and sub-samples of the data, in order to explore how measurement of the outcome may affect the conclusions. We evaluate the fit of these specifications using two different measures: the “percent correctly predicted”, and the “area under the receiver operating characteristic curve” (AUC)¹⁰. The two measures convey similar information, with the difference that the “percent correctly predicted” measure does not use information about the error term distribution, while the AUC does. The AUC is our preferred measure as it incorporates this additional information on the error distribution, which is assumed to be standard normal in our case.

¹⁰The area under the curve can be interpreted as the percent of time the empirical model would classify an instance of sustained poverty reduction and an instance without sustained poverty reduction correctly, compared to a random guess, which would get it right 50 percent of the time.

Column 1 reports our baseline specification. The variables included in this column are the ones dictated by the model. The fixed cost parameter is positive and statistically significant: $\gamma_1 = 0.70$ (s.e. = 0.21), or 700 million people, consistent with our premise of increasing returns to scale.

The effects of both domestic and international market size on sustained poverty reduction are large and significant. The coefficient on the size of the domestic middle class as a share of the population is $\lambda_1 = 0.50$ (s.e. = 0.21), implying that moving from zero to 10 percent of the population in the global middle class is equivalent of adding 50 million people to the population. Recall that variables are scaled, so that the results can be interpreted in terms of people with average income below that of those belonging to the middle class.

The effects of the international market size are also substantial, both when measured by relative population (in 1,000s of people), where $\lambda_2 = 0.15$ (s.e. = 0.09), and when measured by the relative income per capita, $\lambda_3 = 0.04$ (s.e. = 0.02). To understand the magnitude of these effects, consider the situation of Afghanistan, with a population of approximately 35 million. Suppose Afghanistan contemplates whether to integrate with one of its neighbors, either Pakistan, with population of 200 million, or the Islamic Republic of Iran, with population of 80 million. In terms of population, Iran is 2.3 times larger, and Pakistan is 5.7 times larger than Afghanistan. According to the coefficient estimate, opening up to an integrated market of the same population adds the equivalent of 154,000 people to average market size. The multiple of this would be greater if the country integrated with Pakistan. However, the effect of relative income per capita on market size is more important than relative population in the model. According to our coefficient estimate, joining a market with the same relative income per capita, is equivalent to 40 million people on average in our sample. Pakistan, which has a 3 times greater income per capita than Afghanistan would be worth an additional 120 million additional people. However, Iran has income per capita 10 times larger than Afghanistan, and so integration would yield the equivalent of 400 million more people. In this example, Iran is a much more valuable market when one accounts for population and income. Though there are gains to having a large market in terms of population, the main incremental value comes from trading partners' purchasing power. This suggests that so called "South-South" integration between countries of similar incomes will be less valuable than "North-South" integration between countries of different incomes.

Turning to the components of variable profit, the estimate of the constant, $\beta_1 = 1.79$ (s.e. = 0.69), is positive and statistically significant at standard levels. The coefficient on export growth in the period preceding sustained poverty reduction is positive and statistically significant, $\beta_2 = 5.63$ (s.e. = 2.73), suggesting that exports serve indeed as an income boost. However, contrary to

the predictions of the closed economy version of the model, the effect of agricultural productivity growth, $\beta_3 = -5.24$ (s.e. = 7.60), is negative, though not significant.

One reason for this inconclusive result is that, as discussed earlier, in an open economy setting, the effect of a Hicks-neutral increase in agricultural productivity is ambiguous depending on whether the comparative advantage forces emphasized in the work of Corden and Neary (1982) and Matsuyama (1992) or the income boosting effect of technological change dominate. Our setting includes several economies that are partially open, and hence the estimation confounds these two mechanisms that may be simultaneously at work. To investigate whether the sign of agricultural productivity changes when the sample includes only closed economies, we also estimate in column 3 a specification in which the agricultural productivity variable is interacted with a dummy that takes the value of 1 if a country has not signed any trade agreements at all and is thus “closed” according to our definition (a complete list of “closed economies” is provided in the Online Supplement, Table S1). The sign of the interaction becomes positive (consistent with the theory), but the standard errors of both agricultural productivity variables are very large. Just as in the comparison of means between samples in Table 3, we are not able to establish a statistically significant relationship between agricultural productivity growth and sustained poverty reduction. To examine whether factor bias of technological change in agriculture affects our results, we also estimated specifications that used labor productivity in agriculture (rather than Hicks-neutral technical change) as a control. However, the results remained noisy and inconclusive. A possible interpretation of our (non)-results on agricultural productivity is that in an open economy setting, there are many channels through which agricultural productivity growth affects industrialization, with some channels implying a positive, and other channels implying a negative impact. The relevance of these channels likely differs across countries, so that the attempt to estimate a homogeneous effect based on cross-country variation does not produce any statistically significant results.

Having estimated the coefficients of the profit function, and confirmed both the presence of economies of scale and the significance of the market size variables, we now examine what these coefficients imply for our outcome, sustained poverty reduction. In column 1, our baseline specification, the threshold market size to achieve sustained poverty reduction is $\hat{S} = 325$ million people), where those people have purchasing power less than the global middle class. This implies that a large market indeed is required for sustained poverty reduction. This market size can be achieved in a small country however through international trade agreements, or through a more equitable income distribution. The coefficients in parameter vector λ convert domestic and international market size variables into units of population, which allows one to determine what it will take for a

given country to meet the threshold. It is clear therefore that income distribution and international integration will be relatively more important for small countries. Large countries, for instance India and China, have been able to meet this threshold on the basis of population alone.

Regarding the fit of the model, in column 1, $AUC = 0.678$, indicating the model has some predictive power relative to a random guess.¹¹

4.1 Robustness

Column 2 in Table 4 reports results of a specification similar to the baseline specification in column 1, with one difference: instead of using the relative income per capita as a measure of international integrated market size, we use the global middle class share of the integrated market. The results remain strong and in line with the theoretical model, but are noisier than in column 1.¹²

The comparison of the two specifications is interesting in its own right as it provides further (indirect) evidence in support of the mechanisms postulated in the theory. The use of the middle class share of the integrated market instead of relative income per capita affects the magnitudes of two coefficients in a meaningful way: the coefficient on the domestic market size is roughly a third of the coefficient in column 1, and the coefficient on the fixed cost variable is about half the size of the corresponding coefficient in column 1. As a result of the latter, the estimated market size threshold decreases from 325 million people in column 1 to 179 million people in column 2. This decrease is intuitive when interpreted through the lens of the model. Implicit in the use of the integrated market middle class share is the assumption that both domestic and foreign (integrated market) members of the upper class bear the fixed cost of the IRS technology. But if foreign capital can be used to fund the fixed cost of the IRS technology, then the domestic market size constraint becomes less binding as reflected in the decrease of the domestic market threshold. In contrast, when relative income per capita is used, the implicit assumption is that only domestic residents belonging to the upper class fund the fixed cost, while foreign residents (belonging to both the middle and upper classes) contribute to profits. In this case, the threshold domestic market size needs to be larger to support the covering of the fixed cost.

In column 4 of Table 4, we estimate the model using an alternative definition of international integration, according to which only trade agreement provisions covering trade in goods and services

¹¹To provide another reference point for this number, we note that Kleinberg et al. (2017) develop a machine learning tool that can improve on a human judge's decisions to offer or deny bail in New York City. It has an AUC of 0.707, which they consider a good fit.

¹²This is likely due to the fact that the Gini coefficient required to construct the integrated market middle class share is missing for several country-year pairs, so that we use the average (across years) Gini coefficient for each country instead. As a result, the integrated middle class variable is a noisier measure of purchasing power than relative income per capita (see the Online Supplement for details).

are included (so we exclude provisions covering labor and capital flows). This definition is more in line with the spirit of the model that highlights the role of sales and market size in development. Nevertheless, our preferred measure is the one covering all provisions related, directly or indirectly, to trade, as typically, provisions in trade agreements are negotiated jointly. The results in column 4 are almost identical to those in column 1.

In column 5, we rerun the estimation on a select subsample of the data, in order to test whether our results are affected by the additional volatility introduced from the use of income surveys to measure poverty. Specification 5 includes only consumption surveys, reducing the sample to 238 observations. The results are similar as in specification 1, but the standard errors are higher.

In column 6, we consider alternative five-year windows over which the evolution of poverty is examined, in order to explore the sensitivity of our results to the windows selected. When calculating sustained poverty reduction instead of the windows 1981-85, 1986-90, etc. we shift the window one year back to be 1980-84, 1985-89, etc. Right hand side variables are again matched on the first year of the window. Most coefficients in this specification are qualitatively similar to the results in column 1, with the exception of the one on agricultural productivity growth, which now becomes positive, though it remains insignificant. The market size threshold also increases. Based on these results, we conclude that our qualitative results, suggesting large thresholds market size thresholds and an important role of market size for development, are robust to the windows selected, but the exact magnitude of the effects may be affected by the choice of the window.

In the Appendix we report additional robustness checks. Column 1 in Table A1 repeats the results of our baseline specification (column 1 in Table 4) to facilitate comparisons. Column 2 of Table A1 reports results of a specification that includes only geographic and institutional controls. We include variables that the literature has shown to be highly correlated with economic development: the share of land in a tropical climate, the share of land in a desert climate, distance to coast, ruggedness, British legal origins, and French legal origins (the omitted category is socialist legal origins).¹³ These variables have high explanatory power as expected: the simple Probit that includes all these variables has a slightly better fit than our baseline specification (AUC: 0.697 relative to an AUC of 0.678 in our baseline model). However, from a policy point of view, this specification is not particularly useful as these geographic and institutional controls cannot be targeted by policy. The fact that our demand-side-framework inspired specification does almost as well as this specification in terms of fit is reassuring.

¹³The data for the share of land in tropical climate, the share of land in desert climate, average distance to ice-free coast and terrain ruggedness are from Nunn and Puga (2012). For origins of legal system (i.e., English, French or socialist) we use data from La Porta et al. (1999).

In column 3 of Table A1, we consider an alternative definition of sustained poverty reduction that smooths out short term increases in the poverty rate. This alternative definition is equal to one in instances of sustained poverty reduction as defined by our baseline measure, but is also equal to one if the poverty rate has fallen by at least 5 percent between the first and last years of the five-year window. The market size variables are similar to those in column 1, but the threshold market size falls to 126 million people, as expected because the alternative measure is less demanding in classifying instances of sustained poverty reduction. This result suggests poverty reduction can be achieved with a smaller market size, though that it may not be sustained in all periods, as our baseline outcome variable requires.

Column 4 in Table A1 examines the implications of using higher frequency household surveys. This specification includes only surveys in which 0, 1 or 2 years of poverty headcounts are observed, reducing the sample to 200 observations. The signs and statistical significance of the estimated coefficients are robust to using this smaller sample, but the magnitudes change. The main takeaway from these results is that the inclusion of surveys with many missing observations for poverty does not push the estimates in a qualitatively different direction.

Columns 5 and 6 in Table A1 report results from specifications that have a weaker link to the theoretical framework. In column 5, the domestic middle class variable is lagged by 5 years to address potential concerns about reverse causation. Note that the middle class in our baseline specification is measured at the onset of each poverty-reduction-window, so by construction, it is not contemporaneous to poverty change. By using its lag (i.e., measuring it at the onset of the previous 5-year window), we lose 5 years for each country, so that we end up with 285 observations. The results from this specification are similar to those in column 4, in the sense that the coefficient signs remain robust, but the magnitudes change. The standard errors of the key variables become slightly larger due to the lower number of observations. In column 6, we include in addition to the variables suggested by the theoretical framework, initial (in 1981, the first year in our sample) log GDP per capita, to see if there is evidence of conditional convergence in the data. The coefficient on log GDP per capita is insignificant. The remaining results become – if anything - stronger than before, in the sense of being even more supportive of the mechanisms postulated by the theory, but the magnitudes are too large to be plausible and are hard to interpret given that the theoretical model has no implications for (conditional) convergence or divergence. In general, the results in specifications 5 and 6 are qualitatively similar to those in our baseline specification, but given that both specifications have a weaker link to the theoretical framework, we do not attempt to draw any quantitative conclusions from them.

In Table A2, we depart even further from the theory by estimating specifications with a continuous dependent variable: the percent change in the poverty rate. Note that our model has nothing to say about the magnitude of poverty reduction. Accordingly, we cannot use the results from these specifications to estimate market size thresholds. However, the continuous specifications allow us to use more information (i.e., information on the poverty reduction rate every year) and to examine whether the associations between poverty reduction and the key variables in our model are robust. Further, we experiment with various sets of fixed effects, which are useful in highlighting which variation in the data identifies our key parameters. In all specifications in Table A2, we do not use periods in which countries have eliminated extreme poverty, as this would artificially blow up their poverty reduction rate in that period (the denominator would be close to zero). The results in column 1 of Table A2 are the closest analog to those reported in our baseline specification, and the coefficients appear robust in terms of signs and statistical significance. Note that in addition to the variables we use in Table 4, we also include the total domestic population on the right-hand side (in Table 4, we use a normalization that sets this coefficient equal to 1, and which allows us to translate units of market demand to units of population). The estimated coefficient on population in column 1 of Table A2 is positive and statistically significant. Column 2 adds time fixed effects. The addition of time effects renders our two measures of integrated market size (relative population and relative income per capita) insignificant, showing that the variation in these two measures comes primarily from time variation in the signing of trade agreements. Column 3 omits time fixed effects, but uses country fixed effects instead. In this case, it is the total domestic population that changes sign and becomes insignificant, suggesting that the variation in domestic population is identified based on cross-country variation (note however, that this is not the case for domestic middle class). Given that the geographic and institutional control used in Table A1 are constant over time, their role is similar to that of fixed effects. Therefore, in column 4, we also experimented with a specification that omits the country fixed effects, but uses the full set of geographic and institutional controls instead. This makes the coefficient on domestic population positive again, but it remains insignificant.

As noted at the outset, empirical implementation of the theoretical framework requires several judgment calls. The various robustness checks suggest that our main results are robust, at least in qualitative terms, to these judgment calls, though the magnitude of the effects may be less so. Overall, the results support the hypothesis that that poverty reduction is associated with a large market size.

5 Counterfactuals

To evaluate the effect of international integration and the income distribution on sustained poverty reduction, we simulate several counterfactual economies and compare them to the status quo (with current levels of international integration and equality). The first is a counterfactual closed economy without international integration, in which $\lambda_2 = \lambda_3 = 0$. This could be understood as the development policy doomsday scenario, in which comparative advantage from trade becomes irrelevant for sustained poverty reduction. The second is a counterfactual maximum equality economy, in which the share of the population in the middle class is calculated using Equation 10 and the current national GDP per capita, but a Gini coefficient of 0.20, the value for Slovakia in 1992, the lowest in the sample. This could be understood as a scenario where national income stays the same, but government redistributes income to about the extent achieved in the Czechoslovak Socialist Republic, which dissolved in 1989. The third is a fully integrated counterfactual, in which the relative population and income of the integrated market are equal to the relative population and income of the whole world. This could be understood as a maximal trade liberalization scenario in which the country signs deep trade agreements establishing and protecting economic integration rights with every country. Table S2 in the Online Supplement reports estimates of the status quo market size for each country period in our data set, as well as each counterfactual market size. All these scenarios should be thought of as polar cases that serve the purpose of helping us quantify the contributions of equality and international market size towards poverty reduction, rather than realistic policies on the table. We note however that in principle, our framework with the detailed information on specific provisions of trade agreements it provides, could be used to also assess the effects of specific liberalization measures (e.g., signing specific provisions of trade agreements with particular countries).

We summarize these results in two figures. Figure 5 shows the estimated status quo market size, calculated as $M_{it}\hat{\lambda}$, and the three counterfactual market sizes for select economies, using data from 2011, the first year of the last window in our sample, to provide a recent view. The blue set of columns correspond to the status quo, in which market size is calculated using the λ reported in column 2 of Table 4. The orange columns show a market size estimate in which $\lambda_2 = \lambda_3 = 0$, so market size is determined only by population and the size of the middle class. The khaki columns show the maximum equality scenario and the dark green columns show the fully integrated scenario.

The five economies in Figure 5, China, Côte d’Ivoire, Ghana, Indonesia, and Philippines have large enough markets to support sustained poverty reduction under the status quo. Three of them, Côte d’Ivoire, Ghana, and Philippines would not be large enough to sustain poverty reduction as a

closed economy, but the others would, in part due to their larger populations. For these economies, fully integrating produces some gains, but because these economies already have substantial domestic markets, the relative value of the international market is less. In all countries except Côte d’Ivoire, the maximum equality counterfactual increases market size, because their GDP per capita is high enough that increasing equality increases the share of the population making between \$11-\$110 per day. In contrast, in Côte d’Ivoire, equality stands in the way of industrialization, because the economy is too poor to support the increasing returns sector.

Figure 6 shows average market sizes averaged over deciles of GDP per capita, weighting observations by population to provide a poverty reduction relevant view of the heterogeneity across the national income distribution. Notably, in the closed economy scenario, it is not until the sixth decile of GDP per capita, that the market becomes large enough to meet the estimated threshold $\hat{S} = 325$ under a closed economy. International integration appears to help. In all deciles, the status quo market size is on average greater than the threshold. For the lowest income countries, the market size in the status quo scenario is on average greater than the threshold, but barely so. This suggests that, if the value of international markets remains as it has in the past, most countries should be able to achieve sustained poverty reduction. The average market size of the open economy at present however does not go far above the threshold.

6 Concluding Remarks

We started this project in a quest for a new vision for development in an era of rising economic nationalism and increasing automation. We come away with renewed appreciation of what international integration has done in the past for developing countries, especially those with small populations. At the same time, the way forward is more uncertain than ever.

Waning multilateralism makes deep economic integration, especially with richer countries, via additional provisions, for instance regarding the mobility of labor (including the unskilled and professional tradespeople) unlikely. Regionalism may provide an appealing alternative; in fact, many countries in Asia and Africa have embraced a regionally-focused approach in recent years, as evidenced by regional trade agreements, such as RCEP and AfCFTA. However, the message of our analysis is that such integration will be less valuable if it does not involve richer countries. Recent demands for “friendshoring”, i.e., trading primarily with a country’s “friends,” or political allies, are likely to induce high-income countries to trade more with each other and less with developing countries, as P. Goldberg and Reed (2023) show. Whether such demands are justifiable based on geopolitical, environmental, labor, or health concerns is an open question. But they do imply a cost

for developing countries in the form of higher poverty rates, if such countries end up cut off from lucrative foreign markets due to them having lax environmental or labor standards or because of security concerns.

In such an environment, redistribution targeting the poor *and* the middle class becomes more important than ever. While direct aid to the poor is a valuable tool to alleviate poverty, resources must also be made available to broaden the middle class, who sustain the value of the market. The presence of imperfect competition underlying our model suggests that redistribution of firm profits in particular is important. Efforts to assist households in accumulating equity shares may therefore be especially useful policies for redistribution, as opposed to the redistribution of wage income. Along the same lines, wealth redistribution, for example in the form of major land reforms, may be a development promoting strategy in countries characterized by large inequalities.

A sizeable middle class is even more important if the future of export growth is in services, as for instance Baldwin (2020) argues. Services are a highly heterogeneous category, but certain components of services, business and IT services in particular, could in principle play the same role that manufacturing has played in the past in promoting development. Yet the service sector remains highly protected, even within highly integrated areas such as the European Union and the United States, and, not surprisingly, international trade in services remains limited despite having grown fast in recent years (Mattoo, 2018). This state of affairs suggests that services have the potential to become the new frontier of globalization, if trade restrictions were lifted and country-specific regulations prohibiting trade were harmonized across countries. However, such policy changes seem highly unlikely in the present environment, especially since trade in services might threaten many jobs in advanced economies that have so far been sheltered from direct competition from low-wage countries. Against this background, growth in the service sector will require a large domestic population and middle class that can support the fixed costs of new businesses and technologies, generate profits, and set off the virtuous cycle of development postulated in our framework. The recent experience of India, the most populous country of the world as of April 2023, is consistent with this hypothesis.

The stylized empirical model presented abstracts from several other factors that may influence the path of development, even within a demand-side framework. For instance, we consider one poor country at a time, abstracting from the possibility of competition across countries. The entry of China into world markets represented such competition for many low-income countries. In our framework, this was captured by the decrease in the relative income of the integrated market. But the effects of China on developing countries goes beyond its impact on relative income. In a world in

which poor countries compete against each other for lucrative, high-income markets, international integration may be a less promising path to development.

The prospect of competition across countries as well as the importance of the increasing returns sector for development raise the question of whether there is scope for industrial policy. Industrial policy is back in both the United States and Europe, partly as a place-based policy, partly in response to geopolitical concerns, and partly in an effort to address climate change. While its merits are hotly debated, there has been less interest in the role of industrial policy in developing economies. Our results suggest a potential motivation for analyses of industrial policy from a development perspective. The central idea in our paper is that the opportunity to build an industry that helps lift people out of extreme poverty is limited by scale economies. Our approach for identifying the importance of these scale economies and hence market size constraints is to estimate the threshold market size needed to transition out of extreme poverty. But we do not engage with the intermediate steps, with the process of this transition; the specifics, e.g., which industries are characterized by scale economies, does it take one or multiple industries, one or multiple firms, the role of vertical integration, etc., are in the background. If it is specific industries or firms that drive the transition process, how can these be identified? And should they be supported through industrial policy?

These questions are outside the scope of the present work, but we hope that future research will take them on.


References

- Acemoglu, Daron, Simon Johnson, and James A Robinson (2001). “The colonial origins of comparative development: An empirical investigation”. In: *American Economic Review* 91.5, pp. 1369–1401.
- Agarwal, Ruchir and Tristan Reed (2022). “Financing Vaccine Equity”. In: *Oxford Review of Economic Policy* 38.4, pp. 833–850.
- Alfaro-Urena, Alonso, Isabela Manelici, and Jose P Vasquez (2020). *The Effects of Joining Multi-national Supply Chains: New Evidence from Firm-to-Firm Linkages*. Tech. rep.
- Angrist, Noam, Simeon Djankov, Pinelopi K Goldberg, and Harry A Patrinos (2021). “Measuring human capital using global learning data”. In: *Nature* 592.7854, pp. 403–408.
- Atkin, David, Arnaud Costinot, and Masao Fukui (2021). *Globalization and the Ladder of Development: Pushed to the Top or Held at the Bottom?* Tech. rep. National Bureau of Economic Research.
- Atkin, David, Amit K Khandelwal, and Adam Osman (2017). “Exporting and firm performance: Evidence from a randomized experiment”. In: *The Quarterly Journal of Economics* 132.2, pp. 551–615.
- Autor, David H, David Dorn, and Gordon H Hanson (2016). “The China shock: Learning from labor-market adjustment to large changes in trade”. In: *Annual Review of Economics* 8, pp. 205–240.
- Aw, Bee Yan, Mark J. Roberts, and Daniel Yi Xu (June 2011). “RD Investment, Exporting, and Productivity Dynamics”. In: *American Economic Review* 101.4, pp. 1312–44.
- Baldwin, Richard (2008). *Big-think regionalism: a critical survey*. National Bureau of Economic Research Working Paper No. 14056.
- Baldwin, Richard (2020). *The Globotics Upheaval: Globalization, Robotics, and the Future of Work*. Oxford University Press.
- Banerjee, Abhijit and Esther Duflo (2005). “Growth theory through the lens of development economics”. In: *Handbook of Economic Growth* 1, pp. 473–552.
- Banerjee, Abhijit and Esther Duflo (2008). “What is middle class about the middle classes around the world?” In: *Journal of Economic Perspectives* 22.2, pp. 3–28.
- Banerjee, Abhijit, Dean Karlan, and Jonathan Zinman (2015). “Six randomized evaluations of microcredit: Introduction and further steps”. In: *American Economic Journal: Applied Economics* 7.1, pp. 1–21.
- Barro, Robert J (2000). “Inequality and Growth in a Panel of Countries”. In: *Journal of Economic Growth* 5.1, pp. 5–32.
- Barro, Robert J (2008). *Inequality and growth revisited*. Working Paper Series on Regional Economic Integration No. 11. Asian Development Bank.
- Bartelme, Dominick G, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare (2019). *The textbook case for industrial policy: Theory meets data*. Tech. rep. National Bureau of Economic Research.
- Baulch, Bob and John Hoddinott (2000). “Economic mobility and poverty dynamics in developing countries”. In: *The Journal of Development Studies* 36.6, pp. 1–24.
- Birdsall, Nancy, Carol Graham, and Stefano Pettinato (2000). “Stuck In The Tunnel: Is Globalization Muddling The Middle Class?” In.
- Bresnahan, Timothy F and Peter C Reiss (1991). “Entry and competition in concentrated markets”. In: *Journal of Political Economy* 99.5, pp. 977–1009.

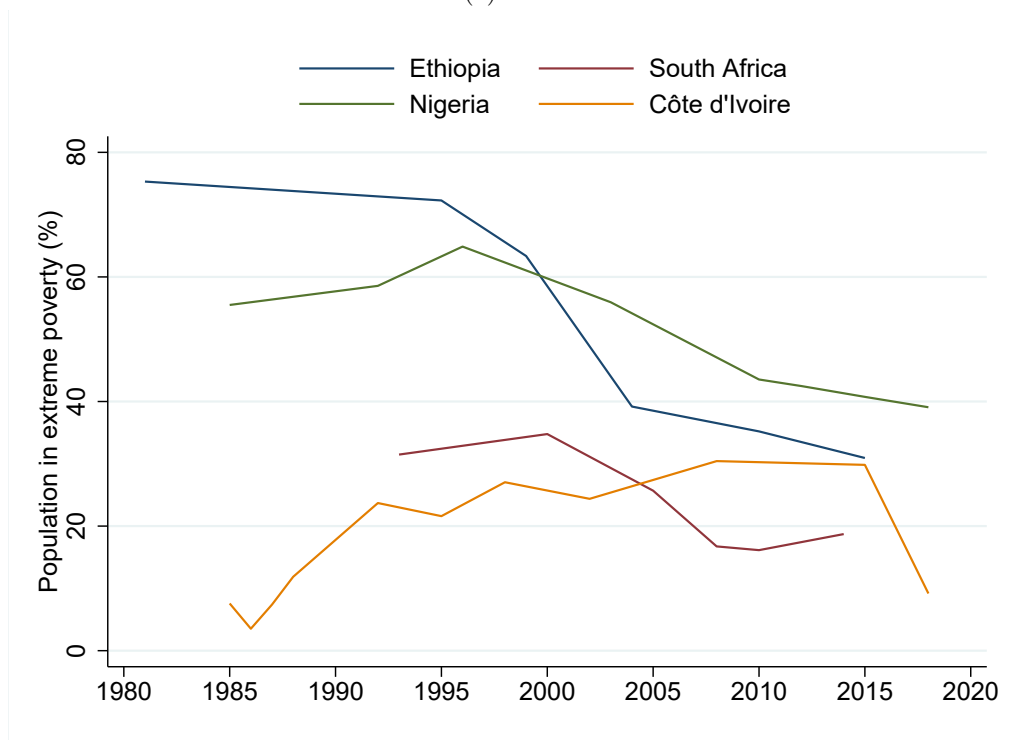
- Bustos, Paula (2011). “Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms”. In: *American Economic Review* 101.1, pp. 304–40.
- Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli (2016). “Agricultural productivity and structural transformation: Evidence from Brazil”. In: *American Economic Review* 106.6, pp. 1320–65.
- Carter, Michael R and Christopher B Barrett (2006). “The economics of poverty traps and persistent poverty: An asset-based approach”. In: *The Journal of Development Studies* 42.2, pp. 178–199.
- Chiquiar, Daniel and Martín Tobal (2019). *Global Value Chains in Mexico: A Historical Perspective*. Working Paper 2019-06. Banco de México.
- Corden, W Max and J Peter Neary (1982). “Booming sector and de-industrialisation in a small open economy”. In: *The Economic Journal* 92.368, pp. 825–848.
- Costinot, Arnaud, Dave Donaldson, Margaret Kyle, and Heidi Williams (2019). “The more we die, the more we sell? a simple test of the home-market effect”. In: *The Quarterly Journal of Economics* 134.2, pp. 843–894.
- Deaton, Angus (2005). “Measuring poverty in a growing world (or measuring growth in a poor world)”. In: *Review of Economics and Statistics* 87.1, pp. 1–19.
- Desai, Raj M and Homi Kharas (2017). *Is a growing middle class good for the poor? Social policy in a time of globalization*. Global Economy and Development Working Paper No. 105. Brookings Institution.
- Easterly, William (2001). “The middle class consensus and economic development”. In: *Journal of Economic Growth* 6.4, pp. 317–335.
- Easterly, William (2019). *In Search of Reforms for Growth: New Stylized Facts on Policy and Growth Outcomes*. National Bureau of Economic Research Working Paper No. 26318.
- Fan, Tianyu, Michael Peters, and Fabrizio Zilibotti (2021). *Service-led or service-biased growth? Equilibrium development accounting across Indian Districts*. Tech. rep. National Bureau of Economic Research.
- Feenstra, Robert C, Robert Inklaar, and Marcel P Timmer (2015). “The next generation of the Penn World Table”. In: *American Economic Review* 105.10, pp. 3150–3182.
- Ferraz, Claudio, Frederico Finan, and Dimitri Szerman (2015). *Procuring firm growth: the effects of government purchases on firm dynamics*. National Bureau of Economic Research Working Paper No. 21219.
- Ferreira, Francisco HG, Phillippe G Leite, and Martin Ravallion (2010). “Poverty reduction without economic growth?: Explaining Brazil’s poverty dynamics, 1985–2004”. In: *Journal of Development Economics* 93.1, pp. 20–36.
- Fuglie, Keith Owen, Sun Ling Wang, V Eldon Ball, et al. (2012). *Productivity growth in agriculture: an international perspective*. CABI.
- Goldberg, Pinelopi and Tristan Reed (2023). *Is the Global Economy Deglobalizing? And if so, why? And what is next?* Brookings Papers on Economic Activity, Spring 2023. Brookings Institution.
- Goldberg, Pinelopi K and Nina Pavcnik (2004). “Trade, inequality, and poverty: What do we know? Evidence from recent trade liberalization episodes in developing countries”. In: *Brookings Trade Forum*, pp. 223–269.
- Gollin, Douglas (2010). “Agricultural productivity and economic growth”. In: *Handbook of Agricultural Economics* 4, pp. 3825–3866.

- Gollin, Douglas (2014). “The Lewis model: A 60-year retrospective”. In: *Journal of Economic Perspectives* 28.3, pp. 71–88.
- Gollin, Douglas, Remi Jedwab, and Dietrich Vollrath (2016). “Urbanization with and without industrialization”. In: *Journal of Economic Growth* 21.1, pp. 35–70.
- Gollin, Douglas, Stephen L Parente, and Richard Rogerson (2004). *Structural Transformation and Economic Development*. Tech. rep. Working paper, Arizona State University.
- Harrison, Ann, ed. (2007). *Globalization and poverty*.
- Hausmann, Ricardo, Lant Pritchett, and Dani Rodrik (2005). “Growth accelerations”. In: *Journal of Economic Growth* 10.4, pp. 303–329.
- Helpman, Elhanan and Paul Krugman (1985). *Market structure and foreign trade: Increasing returns, imperfect competition, and the international economy*. MIT press.
- Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi (2014). “Growth and structural transformation”. In: *Handbook of Economic Growth*. Vol. 2. Elsevier, pp. 855–941.
- Hofmann, Claudia, Alberto Osnago, and Michele Ruta (2017). *Horizontal depth: a new database on the content of preferential trade agreements*. The World Bank.
- Keane, Michael P and Eswar S Prasad (2002). “Inequality, transfers, and growth: New evidence from the economic transition in Poland”. In: *Review of Economics and Statistics* 84.2, pp. 324–341.
- Kharas, Homi (2010). “The emerging middle class in developing countries”. In: *OECD*.
- Kharas, Homi (2017). “The unprecedented expansion of the global middle class: An update”. In.
- Kleinberg, Jon et al. (2017). “Human decisions and machine predictions”. In: *The Quarterly Journal of Economics* 133.1, pp. 237–293.
- Klette, Tor Jakob and Zvi Griliches (1996). “The inconsistency of common scale estimators when output prices are unobserved and endogenous”. In: *Journal of Applied Econometrics* 11.4, pp. 343–361.
- Krugman, Paul (1994). “The fall and rise of development economics”. In: *Rethinking the development experience: Essays provoked by the work of Albert O. Hirschman*, pp. 39–58.
- La Porta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert Vishny (1999). “The quality of government”. In: *The Journal of Law, Economics, and Organization* 15.1, pp. 222–279.
- Leamer, Edward E and James Levinsohn (1995). “International trade theory: the evidence”. In: *Handbook of International Economics* 3, pp. 1339–1394.
- Lileeva, Alla and Daniel Trefler (2010). “Improved Access To Foreign Markets Raises Plant-Level Productivity...For Some Plants”. In: *The Quarterly Journal of Economics* 125.3, pp. 1051–1099.
- Lund, Susan et al. (2019). *Globalization in Transition: The Future of Trade and Value Chains*. Tech. rep. McKinsey Global Institute.
- Matsuyama, Kiminori (1992). “Agricultural productivity, comparative advantage, and economic growth”. In: *Journal of Economic Theory* 58.2, pp. 317–334.
- Mattoo, Aaditya (2018). *Services Globalization in an Age of Insecurity: Rethinking Trade Cooperation*. The World Bank.
- McKenzie, David and Christopher Woodruff (2014). “What are we learning from business training and entrepreneurship evaluations around the developing world?” In: *The World Bank Research Observer* 29.1, pp. 48–82.
- Milanovic, Branko (2013). “All the Ginis dataset”. In: *The World Bank*.

- Murphy, Kevin M, Andrei Shleifer, and Robert Vishny (1989). “Income distribution, market size, and industrialization”. In: *The Quarterly Journal of Economics* 104.3, pp. 537–564.
- Nunn, Nathan and Diego Puga (2012). “Ruggedness: The blessing of bad geography in Africa”. In: *Review of Economics and Statistics* 94.1, pp. 20–36.
- Ostry, Mr Jonathan David, Mr Andrew Berg, and Mr Charalambos G Tsangarides (2014). *Redistribution, inequality, and growth*. Staff Discussion Note No. 2. International Monetary Fund.
- Pinkovskiy, Maxim and Xavier Sala-i-Martin (2014). “Africa is on time”. In: *Journal of Economic Growth* 19.3, pp. 311–338.
- Ravallion, Martin (2003). “The debate on globalization, poverty and inequality: why measurement matters”. In: *International Affairs* 79.4, pp. 739–753.
- Ravallion, Martin (2009). *The developing world’s bulging (but vulnerable)” middle class*”. World Bank Policy Research Working Paper 4816.
- Ravallion, Martin (2012). “Why Don’t We See Poverty Convergence?” In: *American Economic Review* 102.1, pp. 504–523.
- Rodrik, Dani (2016). “Premature deindustrialization”. In: *Journal of Economic Growth* 21.1, pp. 1–33.
- Rodrik, Dani, Margaret McMillan, and Claudia Sepulveda (2013). “Structural change, Fundamentals, and Growth: An Overview”. In: *Institute for Advanced Study* 23.
- Sachs, Jeffrey D (2001). *Tropical underdevelopment*. National Bureau of Economic Research Working Paper No. 8119.
- Sachs, Jeffrey D and Andrew Warner (1995). “Economic reform and the process of global integration”. In: *Brookings Papers on Economic Activity* 1995.1, pp. 1, 1095, 108–118.
- Shepherd, Andrew and Vidya Diwakar (July 2019). *Pathways to sustained poverty reduction*. Tech. rep. Overseas Development Institution.
- Simanis, Erik (2012). “Reality check at the bottom of the pyramid”. In: *Harvard Business Review*.
- Spence, Michael et al. (2008). “The growth report: Strategies for sustained growth and inclusive development”. In: *Commission on Growth and Development Final Report, Washington, DC*.
- Stiglitz, Joseph E (1996). “Some lessons from the East Asian miracle”. In: *The World Bank Research Observer* 11.2, pp. 151–177.
- Topalova, Petia (2010). “Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India”. In: *American Economic Journal: Applied Economics* 2.4, pp. 1–41.
- Wacziarg, Romain and Karen Horn Welch (2008). “Trade liberalization and growth: New evidence”. In: *The World Bank Economic Review* 22.2, pp. 187–231.
- Werker, Eric (2012). *Learning from Double-Digit Growth Experiences*. International Growth Center Rapid Response Note.
- Winters, L Alan, Neil McCulloch, and Andrew McKay (2004). “Trade liberalization and poverty: the evidence so far”. In: *Journal of Economic Literature* 42.1, pp. 72–115.
- Woodruff, Christopher (2018). *Addressing constraints to small and growing businesses*. Tech. rep. International Growth Centre.
- World Bank (2020). *Poverty and Inequality Platform (database)*. URL: www.pip.worldbank.org.



(a) Asia



(b) Africa

Figure 1: Poverty reduction in select countries

Notes: Extreme poverty is living on \leq \$1.90/day PPP in 2011 US dollars.

Periods of sustained poverty reduction (%)




Figure 2: Sustained poverty reduction

Notes: The sample includes 377 5-year periods, from 94 distinct countries, excluding advanced economies (i.e., those with less than 3% of the population in extreme poverty for all periods in the sample). Extreme poverty reduction is a continuous reduction in the share of the population living on \leq \$1.90/day PPP in 2011 US dollars, assuming a linear trend in the poverty headcount ratio between years of survey data.




Figure 3: The China shock to the World Trade Organization

Notes: Income and population are summed over all member countries for each year.

Relative size of integrated international market

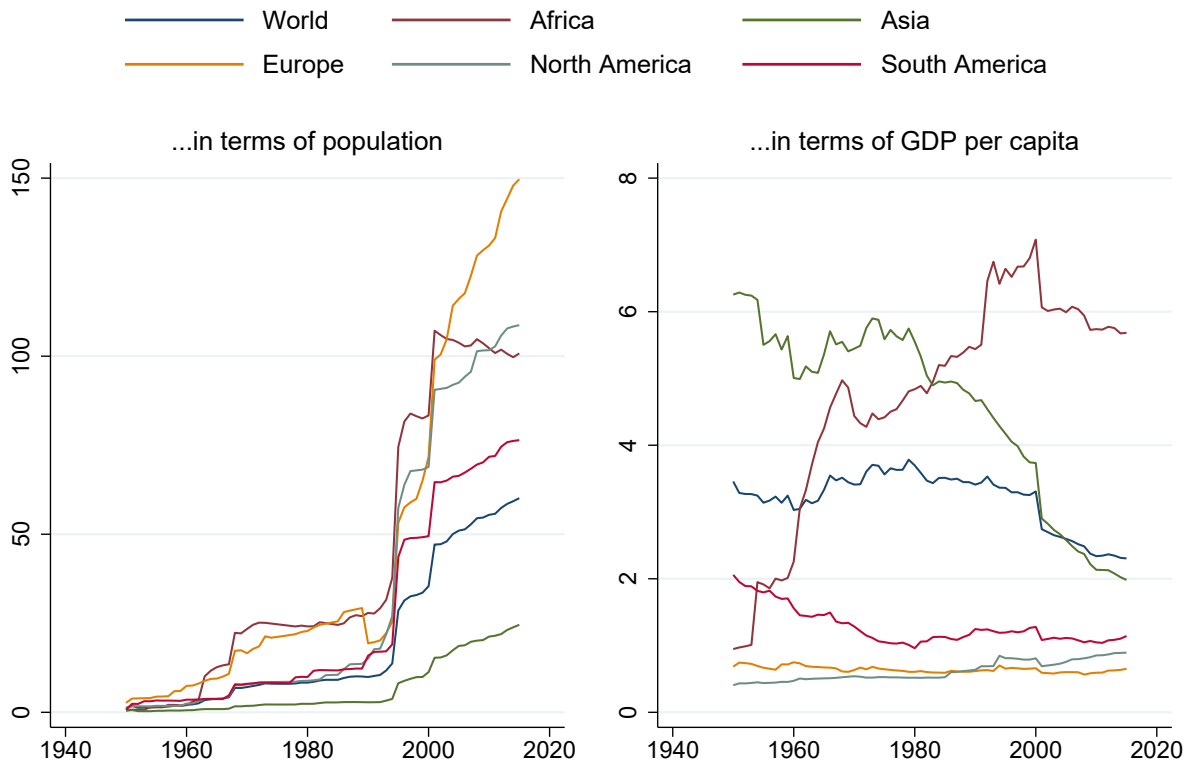


Figure 4: The rise (and fall) in the value of the integrated international market

Notes: Regional average values are weighted by population. If an economy has not signed any trade agreements, relative population and income are set equal to zero.




Figure 5: Counterfactual market sizes in select countries

Notes: Data are for 2011. Market size is calculated using coefficients estimated in column 4 of Table 4. The status quo indicates the market size observed in the data. The closed economy counterfactual is the market size when setting the relative population and income of the integrated market to zero. The maximum equality counterfactual is the market size when the share of the population in the middle class is calculated with the current national GDP per capita, but a Gini coefficient of 0.20, the value for Slovakia in 1992, the lowest in the sample. The fully integrated counterfactual sets the relative population and income of the integrated market equal to the relative income of the whole world, as if the country signed all trade agreements with all countries.

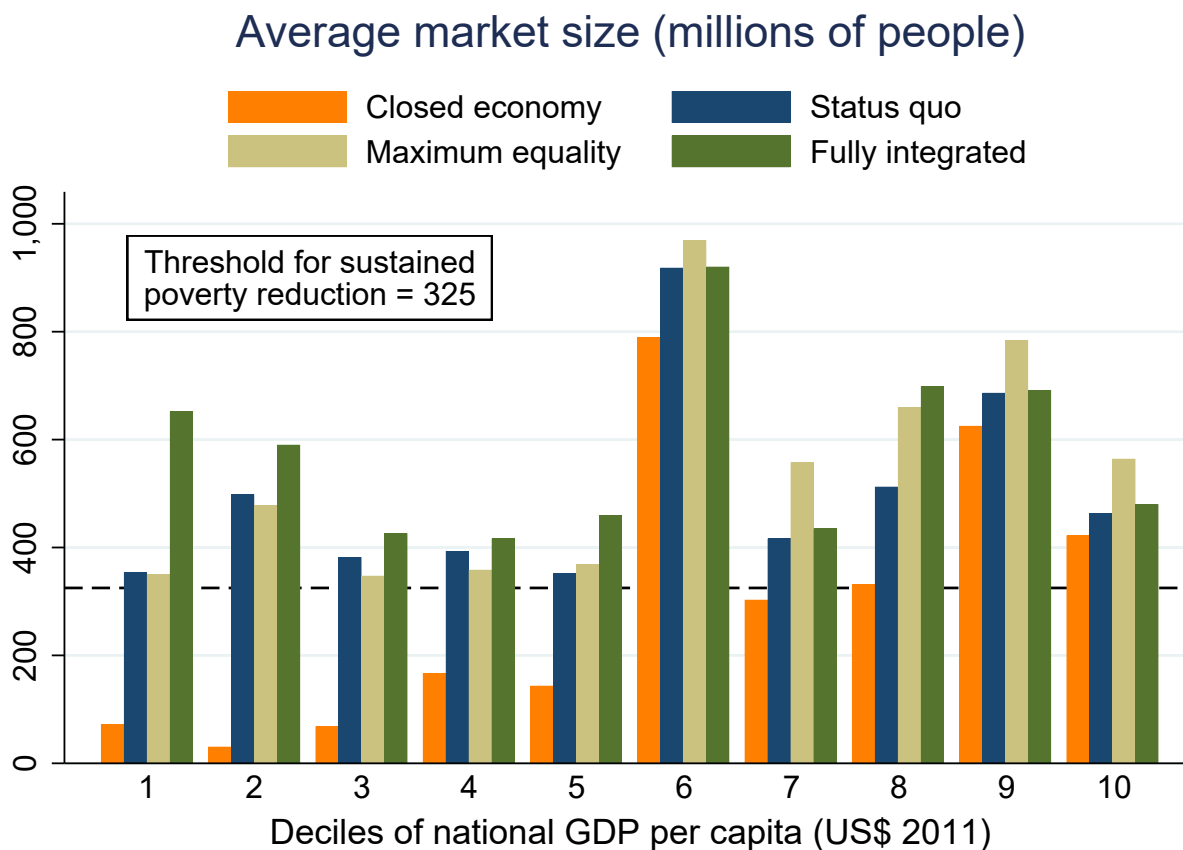


Figure 6: Counterfactual market sizes in full sample, by national GDP per capita

Notes: Averages are weighted by population in the base year. Market size is in units of people outside the middle class, i.e., those consuming less than \$11 per day, PPP in 2011 US dollars. Data are for 2011. Market size is calculated using coefficients estimated in column 4 of Table 4. The status quo indicates the market size observed in the data. The closed economy counterfactual is the market size when setting the relative population and income of the integrated market to zero. The maximum equality counterfactual is the market size when the share of the population in the middle class is calculated with the current national GDP per capita, but a Gini coefficient of 0.20, the value for Slovakia in 1992, the lowest in the sample. The fully integrated counterfactual sets the relative population and income of the integrated market equal to the relative income of the whole world, as if the country signed all trade agreements with all countries.

Table 1: Periods of sustained poverty reduction and growth

		Sustained real per capita GDP growth		
		No	Yes	Total
Sustained extreme poverty reduction	No	79 48%	84 52%	163 100%
	Yes	51 24%	163 76%	214 100%
Total		130 34%	247 66%	377 100%

Notes: Percentages report shares of row totals.
An observation is a country five-year period,
ranging from 1981-2015.

Sources: World Bank Poverty and Inequality
Platform. Penn World Tables 9.1.

Table 2: Core provisions of multilateral trade agreements

Policy Area	General Agreement on Tariffs and Trade (GATT)	World Trade Organization (WTO)	Agreement on Government Procurement (GPA)	Preferential Trade Agreements (PTAs)
<i>A) Establishing basic economic integration rights</i>				
goods	<ul style="list-style-type: none"> • Industrial tariffs • Agricultural tariffs 			<ul style="list-style-type: none"> • Industrial tariffs (WTO+) • Agricultural tariffs (WTO+) • Export taxes (WTO+)
services		<ul style="list-style-type: none"> • General Agreement on Trade in Services (GATS) 		<ul style="list-style-type: none"> • GATS (WTO+)
capital		<ul style="list-style-type: none"> • Agreement on Trade-Related Investment Measures (TRIMS) 		<ul style="list-style-type: none"> • TRIMS (WTO+) • Local content (WTO-X) • Repatriation of capital (WTO-X)
labor				<ul style="list-style-type: none"> • Visa and asylum (WTO-X)
ideas		<ul style="list-style-type: none"> • Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS) 		<ul style="list-style-type: none"> • TRIPS (WTO+) • Intellectual property rights (WTO-X)
<i>B) Protecting these rights (by limiting government discretion to undo them)</i>				
	<ul style="list-style-type: none"> • Customs administration • Anti-dumping (GATT Article VI) • Countervailing measures (GATT Article VI) 	<ul style="list-style-type: none"> • Agreement on the Application of Sanitary and Phytosanitary (SPS) Measures • Technical Barriers to Trade (TBT) Agreement • Agreement on Subsidies and Countervailing Measures (ASCM) • State trading enterprises (GATT Article XVII) 	<ul style="list-style-type: none"> • Public procurement 	<ul style="list-style-type: none"> • Customs administration (WTO+) • Anti-dumping (WTO+) • Countervailing measures (WTO+) • SPS (WTO+) • TBT (WTO+) • Subsidies (WTO+) • State trading enterprises (WTO+) • Public procurement (WTO+) • Competition policy (WTO-X)
<i>Provisions</i>	5	7	1	19
			TOTAL	32

Notes: For preferential trade agreements, WTO+ indicates provisions that are within the scope of the WTO's jurisdiction, whereas provisions indicated by WTO-X are not.

Sources: WTO and Hofmann, Osnago, and Ruta (2017). In Hofmann et. al. provisions related to local content are labeled "investment," and provisions related to repatriation of capital are called "movement of capital".

Table 3: Sample market descriptive statistics

Variables	Mean	Standard Deviation	Min	Max	Differences in means, by sustained poverty reduction status			Standard error
					Mean, Sustained Poverty Reduction = 0	Mean, Sustained Poverty Reduction = 1	Difference	
Domestic market								
Population (billions of people)	0.07	0.20	0.00	1.37	0.04	0.09	0.06	(0.021)
Middle class (share of population)	0.37	0.29	0.00	0.97	0.34	0.38	0.04	(0.03)
Integrated international market								
Relative population (1,000s of people)	0.25	0.48	0.00	4.79	0.19	0.29	0.09	(0.05)
Relative GDP per capita	3.54	3.60	0.00	20.27	3.24	3.76	0.52	(0.374)
Middle class (share of population)	0.41	0.16	0.00	0.83	0.40	0.42	0.03	(0.017)
Past income boosts								
Exports (annual growth)	0.08	0.11	-0.29	0.68	0.06	0.10	0.04	(0.011)
Agricultural TFP (annual growth)	0.01	0.02	-0.08	0.11	0.01	0.01	-0.001	(0.003)

Notes: Sustained poverty reduction is a continuous reduction in the share of the population in extreme poverty (i.e., earning less than \$1.90/day PPP in 2011 US dollars) over a 5 year period, assuming a linear trend between years of survey data. The sample includes 377 observations of 94 countries between 1981-2015, and excludes advanced economies (i.e., those with less than 3% of the population in extreme poverty for all periods in the sample). Middle class is the share of the population earning \$11-110/day PPP 2011, calculated as a function of the average Gini coefficient and real GDP per capita, assuming a log-normal income distribution. Population and GDP of integrated international market is calculated summing all the countries in the world, weighted by the depth of trade agreements signed between them.

Sources: World Bank Poverty and Inequality Platform, Penn World Tables 9.1, World Trade Organization, Hofmann, Osnago, and Ruta (2017), USDA ERS

Table 4: Threshold profit functions

DEPENDENT VARIABLE		(1)	(2)	(3)	(4)	(5)	(6)
		Sustained poverty reduction	Sustained poverty reduction	Sustained poverty reduction	Sustained poverty reduction	Sustained poverty reduction	Sustained poverty reduction
Market size (S)	Middle class share of domestic population	0.503 (0.207)	0.150 (0.117)	0.650 (0.346)	0.504 (0.208)	0.735 (0.488)	0.849 (0.346)
	Relative population of integrated market (1,000s of people)	0.154 (0.089)	0.154 (0.095)	0.183 (0.126)	0.138 (0.08)	0.101 (0.072)	0.048 (0.08)
	Relative GDP per capita of integrated market	0.043 (0.02)		0.056 (0.032)	0.043 (0.02)	0.032 (0.022)	0.078 (0.032)
	Middle class share of integrated market		0.32 (0.222)				
Variable profit (V)	Constant	1.786 (0.686)	1.773 (0.804)	1.480 (0.785)	1.785 (0.686)	3.323 (2.254)	1.094 (0.465)
	Exports (% annual growth)	5.626 (2.73)	7.435 (4.556)	4.492 (2.619)	5.631 (2.733)	7.537 (5.258)	1.994 (1.334)
	Agricultural total factor productivity (annual growth)	-5.239 (7.598)	-14.012 (13.569)	-4.838 (6.043)	-5.232 (7.593)	0.048 (8.962)	5.555 (4.442)
	Agricultural total factor productivity (annual growth) X (Closed == 1)			14.758 (18.095)			
Fixed cost (F)	Constant	0.703 (0.207)	0.389 (0.186)	0.734 (0.208)	0.703 (0.207)	1.138 (0.331)	0.715 (0.207)
Specification		Baseline	Alternative measure of integrated international market	Interaction of agricultural productivity with indicator for closed economy	Integration defined only by goods and service agreements	Consumption survey countries only	Five year windows shifted one year back
Log likelihood		-237.9	-243.7	-237.7	-237.9	-128.3	-240.2
Area under the receiver operating characteristic curve (AUC)		0.678	0.644	0.682	0.678	0.756	0.642
Percent correctly predicted		0.615	0.607	0.623	0.618	0.706	0.576
Threshold market size (billions of people outside middle class)		0.325	0.179	0.166	0.325	0.289	0.545
Number of observations		377	377	377	377	238	368

Notes: Asymptotic standard errors in parenthesis.

APPENDIX: Alternative empirical specifications

Table A1: Alternative empirical specifications

DEPENDENT VARIABLE		(1)	(2)	(3)	(4)	(5)	(6)
		Sustained poverty reduction	Sustained poverty reduction	Sustained poverty reduction	Sustained poverty reduction	Sustained poverty reduction	Sustained poverty reduction
Market size (S)	Middle class share of domestic population	0.503 (0.207)		0.435 (0.207)	0.214 (0.114)		0.524 (0.21)
	Middle class share of domestic population (lagged 5 years)					0.181 (0.15)	
	Log of GDP per capita (1981)						0.044 (0.626)
	Relative population of integrated international market (1,000s of people)	0.154 (0.089)		0.157 (0.089)	0.077 (0.039)	0.078 (0.062)	0.188 (0.093)
	Relative GDP per capita of integrated market	0.043 (0.02)		0.028 (0.02)	0.009 (0.006)	0.018 (0.016)	0.179 (0.019)
Variable profit (V)	Constant	1.786 (0.686)		1.681 (0.686)	7.858 (4.301)	3.272 (2.456)	1.923 (0.675)
	Exports (% annual growth)	5.626 (2.73)		8.065 (2.73)	22.427 (11.392)	8.137 (7.478)	1.517 (3.878)
	Agricultural total factor productivity (annual growth)	-5.239 (7.598)		-12.712 (7.598)	-16.808 (30.613)	-12.129 (22.019)	-0.994 (3.061)
Fixed cost (F)	Constant	0.703 (0.207)		0.272 (0.207)	0.786 (0.306)	0.407 (0.259)	3.183 (8.198)
Other variables	Tropical climate (share of land area)		-0.671 (0.2)				
	Desert climate (share of land area)		1.266 (0.736)				
	Distance to ice-free coast (1,000s of km)		0.297 (0.131)				
	Ruggedness		-0.048 (0.049)				
	British legal origin		0.857 (0.196)				
	French legal origin		0.293 (0.184)				
Specification		Baseline	Probit with only geographic and institutional variables	Alternate poverty reduction measure	Less than two years of poverty observed in window	Middle class share of domestic population lagged 5 years	Control for GDP per capita
Log likelihood		-237.9	-236.2	-225.3	-108.6	-181.8	-237.2
Area under the receiver operating characteristic curve (AUC)		0.678	0.697	0.669	0.750	0.643	0.680
Percent correctly predicted		0.615	0.658	0.634	0.695	0.614	0.629
Threshold market size (billions of people outside middle class)		0.325		0.126	0.084	0.107	1.567
Number of observations		377	377	377	200	285	377

Notes: Asymptotic standard errors in parenthesis.

Table A2: Alternative empirical specifications with a continuous dependent variable

	(1)	(2)	(3)	(4)
VARIABLES	Percent change in poverty rate	Percent change in poverty rate	Percent change in poverty rate	Percent change in poverty rate
Population (billions)	0.135 (0.072)	0.196 (0.091)	-0.092 (0.462)	0.032 (0.067)
Middle class share of domestic population	0.288 (0.108)	0.196 (0.093)	0.609 (0.237)	0.366 (0.120)
Relative population of integrated international market (1,000s of people)	0.070 (0.027)	0.012 (0.024)	0.133 (0.074)	0.049 (0.027)
Relative GDP per capita of integrated market	0.005 (0.006)	-0.004 (0.004)	0.029 (0.024)	0.014 (0.008)
Exports (% annual growth, last 5 years)	0.571 (0.163)	0.280 (0.163)	0.482 (0.183)	0.498 (0.162)
Agricultural total factor productivity (% annual growth, last 5 years)	0.292 (0.749)	-0.087 (0.699)	0.497 (0.936)	0.185 (0.715)
Tropical climate (% of land area)				0.047 (0.064)
Desert climate (% of land area)				0.528 (0.185)
Distance to ice-free coast (1,000s of km)				-0.079 (0.069)
Ruggedness				0.041 (0.018)
British legal origin				-0.202 (0.069)
French legal origin				-0.314 (0.081)
Constant	-0.053 (0.077)	0.052 (0.051)	-0.242 (0.158)	0.085 (0.118)
R-squared	0.058	0.138	0.305	0.127
Number of observations	324	324	324	324
Time fixed effects	No	Yes	No	No
Country fixed effects	No	No	Yes	No

Notes: Asymptotic standard errors in parenthesis. Omits five year periods in which a country has eliminated extreme poverty.

Online Supplement: Data Construction

Sustained Poverty Reduction The headcount series for each country is transformed into a series of periods of sustained poverty reduction in the following four steps. First, for years where there is a missing headcount within a country, a linear trend is estimated between years, and used to interpolate the missing data. Second, the data are segmented into seven mutually exclusive 5-year periods. Third, for all years, we create an indicator for whether the headcount is lower relative to the previous year, referencing either the observed or interpolated value. Fourth, this indicator is used to create an indicator for sustained poverty reduction *throughout* the period. If interpolated and observed poverty has fallen in all years within the five year period, we code the indicator equal to 1 for that period. If the poverty headcount is not declining in every year, we code the variable equal to 0. If an observed or interpolated headcount is not available in all five years, the indicator is missing for that period.

To see how this works, consider the example of Angola, a large high poverty countries. In Angola, the extreme poverty headcount is only observed three times, in 2000, at 36.3 percent, in 2008, at 34.5 percent, and in 2018 at 49.9 percent. Using the linear interpolation, we therefore only observe sustained poverty reduction in each year from 2001 to 2008. When data are segmented into periods, 01-05 is the only five year period for which the change in headcount is observed. We code sustained poverty reduction equal to 1 in this period, because the linear trend is always negative. Since the trend is positive between 2008 and 2018 we code sustained poverty reduction equal to 0 in the subsequent periods 06-10 and 11-15. This example highlights the challenge involved with measuring changes in poverty at the national level; statistical agencies may release numbers infrequently.

International market size Our data on trade agreement provisions come primarily from Hofmann, Osnago, and Ruta (2017), who code the legal content of all provisions of all 279 regional trade agreements in force and notified to the World Trade Organization as of 2015. These authors build on the approach developed by Horn, Mavroidis, and Sapir (2010) for preferential trade agreements involving the United States (US) and European Community (EC). We add to their work by coding provisions linked to three major international (rather than regional) trade agreements, the General Agreement on Tariffs and Trade (GATT), the various agreements of the World Trade Organization (WTO), and the Government Procurement Agreement (GPA).

Table 2 distinguished between two types of provisions. First, there are provisions establishing rights protecting the mobility of goods and services, labor, capital and ideas. Rights over trade in goods, for instance the right to receive the most favored nation tariff, have been enforceable for some time, first under the GATT Article I and then the WTO. Establishment of rights related to services trade remains limited to those areas covered under the General Agreement on Trade in Services (GATS), from which major sectors are excluded, such as maritime services. Provisions enforcing rights over labor and capital mobility are the rarest. Visa and asylum provisions for instance are only protected in regional trade agreements, such as the EC or the Economic Community of West African States (ECOWAS). Rights to free capital flows, such as prohibitions on local content requirements and protections of the right to repatriate profits, have only recently emerged under certain preferential trade agreements. Finally, some provisions enforce rights over the mobility of ideas, via intellectual property rights protection which is controversial, for instance in the pharmaceutical industry where losses in consumer welfare may be substantial (Chaudhuri, Goldberg, and Jia, 2006).

Second, the agreements have additional provisions to protect the specified rights, by limiting

government discretion to undo them. For instance, the Agreement on Subsidies and Countervailing Measures (ASCM) gives rights to the withdrawal of subsidies or the removal of their adverse effects. Countries also have the right, after an investigation, to charge “countervailing” duties on subsidized imports that are found to be hurting domestic producers. There is an argument that, particularly as regards provisions regarding the mobility of ideas, capital and services, trade agreements have become captured by rich-country business elites (Rodrik, 2018) and do not necessarily serve developing countries. Given this, we interpret any observed effect of joining an agreement to be the net of potentially positive and potentially negative effects.

We use the pairwise provisions signed between countries to construct a country-specific measure of the relative size of the international market. As our indicators of relative market size, we use both per-capita income, or per-capita GDP in current US dollars, and population. Formally, for each integrated market M and country i and year t , let

$$\begin{aligned} \text{income}_{it}^M &= \sum_j \rho_{ijt} \text{income}_{jt} \\ \text{population}_{it}^M &= \sum_j \rho_{ijt} \text{population}_{jt} \end{aligned} \tag{7}$$

where ρ_{ijt} equals the number of in force provisions signed between country i and j in year t , as a share of the maximum, 32. Using these statistics, we calculate the relative size of the integrated market, in terms of population and income, for each country year:

$$\begin{aligned} \text{relative population of integrated market}_{it} &= \text{population}_{it}^M / \text{population}_{it} \\ \text{relative income of integrated market}_{it} &= \frac{\text{income}_{it}^M / \text{population}_{it}^M}{\text{income}_{it} / \text{population}_{it}}. \end{aligned} \tag{8}$$

One question is how our legally-determined measure of market integration compares with other measures of liberalization in the literature. Table S1 lists, for 1981, 2001 and 2015, all economies which are considered closed by our measure, given that they have no signatures to any of the treaties in our database. Of current members of the United Nations that existed at the time, in 1981 we count 67 closed economies. Of these, only four, Botswana, Ecuador, Jordan and Thailand were classified as open at the time by Sachs and Warner (1995) and Wacziarg and Welch (2008). Their classifications however are much more likely to classify countries as closed, even when they appear open in our sample. For example, Wacziarg and Welch (2008) consider China and India closed as of 2001, even though India had been a WTO member since 1995 and a member of GATT since 1948. China joined the WTO in that year. According to our treaty-based classification of openness, very few economies remained closed as of 2015, the three largest being Afghanistan, the Islamic Republic of Iran and the Democratic People’s Republic of Korea, and the others being very small states, often with special customs arrangements (i.e., Monaco with France, or the Marshall Islands and Palau with the United States) that may obviate the need for multilateral agreements to obtain market integration. Afghanistan joined the WTO in July 2016.

Income distribution Pinkovskiy and Sala-i-Martin (2009) show that the log-normal distribution provides a good fit to the income distribution in most countries, delivering distributions very similar to those obtained from kernel density estimates, and of superior fit to the gamma and Weibull distributions, two alternatives that also have two parameters.

The middle class share variable combines measures of both equality and average income. To measure equality, we use the Gini coefficient, a standard measure of (in)equality. Since estimation of the Gini coefficient requires a household survey, it is measured infrequently within countries. For this reason, we take the long run average Gini coefficient (G_i) within each country using data from (Milanovic, 2013). For countries in which there are multiple series, we use the series with the most observations. For average income, we use data on real GDP per capita (\bar{Y}_{it}) in 2011 PPP dollars from the Penn World Tables 9.1 (Feenstra, Inklaar, and Timmer, 2015). 2011 PPP dollars are used to match the bounds of the global middle class proposed by Kharas (2010).

Suppose individual daily income y is distributed according to $\ln(y) \sim \mathcal{N}(\mu_{it}, \sigma_{it}^2)$, so that:

$$\begin{aligned}\mu_{it} &= \ln(\bar{Y}_{it}/365) - \frac{\sigma_{it}^2}{2}, \\ \sigma_{it} &= \sqrt{2}\Phi^{-1}\left(\frac{G_i + 1}{2}\right)\end{aligned}\tag{9}$$

where $\Phi^{-1}(\cdot)$ is the inverse normal cumulative distribution function. Aitchison and Brown (1957) first showed the link between the Gini and the parameters of the log-normal distribution, the properties of which are reviewed by Crow and Shimizu (1987).

Then

$$\text{middle class share of population}_{it} = \Phi\left(\frac{\ln(110) - \mu_{it}}{\sigma_{it}}\right) - \Phi\left(\frac{\ln(11) - \mu_{it}}{\sigma_{it}}\right).\tag{10}$$

References

- Aitchison, John and James AC Brown (1957). “The lognormal distribution with special reference to its uses in economics”. In.
- Chaudhuri, Shubham, Pinelopi K Goldberg, and Panle Jia (2006). “Estimating the effects of global patent protection in pharmaceuticals: a case study of quinolones in India”. In: *American Economic Review* 96.5, pp. 1477–1514.
- Crow, Edwin L and Kunio Shimizu (1987). *Lognormal distributions*. Marcel Dekker.
- Feenstra, Robert C, Robert Inklaar, and Marcel P Timmer (2015). “The next generation of the Penn World Table”. In: *American Economic Review* 105.10, pp. 3150–3182.
- Hofmann, Claudia, Alberto Osnago, and Michele Ruta (2017). *Horizontal depth: a new database on the content of preferential trade agreements*. The World Bank.
- Horn, Henrik, Petros C Mavroidis, and André Sapir (2010). “Beyond the WTO? An anatomy of EU and US preferential trade agreements”. In: *The World Economy* 33.11, pp. 1565–1588.
- Kharas, Homi (2010). “The emerging middle class in developing countries”. In: *OECD*.
- Milanovic, Branko (2013). “All the Ginis dataset”. In: *The World Bank*.
- Pinkovskiy, Maxim and Xavier Sala-i-Martin (2009). *Parametric estimations of the world distribution of income*. National Bureau of Economic Research Working Paper No. 15433.
- Rodrik, Dani (2018). “What do trade agreements really do?” In: *Journal of Economic Perspectives* 32.2, pp. 73–90.
- Sachs, Jeffrey D and Andrew Warner (1995). “Economic reform and the process of global integration”. In: *Brookings Papers on Economic Activity* 1995.1, pp. 1, 1095, 108–118.
- Wacziarg, Romain and Karen Horn Welch (2008). “Trade liberalization and growth: New evidence”. In: *The World Bank Economic Review* 22.2, pp. 187–231.

Table S1: Closed economies, as determined by lack of signature to any legally-binding multilateral trade agreement

A) Closed Economies in 1981			
Afghanistan ^a	Djibouti ^a	Libya ^a	Sao Tome and Principe
Albania ^a	Ecuador ^b	Maldives ^a	Saudi Arabia ^a
Algeria	Equatorial Guinea ^a	Marshall Islands ^a	Seychelles ^a
Andorra ^a	Eswatini (Swaziland)	Mexico	Solomon Islands ^a
Angola	Ethiopia	Micronesia, Fed. Sts. ^a	Somalia ^a
Bahamas, The ^a	Fiji	Monaco ^a	Sudan ^a
Bahrain ^a	Guinea	Mongolia	Thailand ^b
Bhutan	Guinea-Bissau	Morocco	Tonga ^a
Bolivia	Iran, Islamic Rep.	Mozambique	Tunisia
Botswana ^b	Iraq ^a	Namibia	Tuvalu ^a
Brunei Darussalam ^a	Jordan ^{a,b}	Nauru ^a	United Arab Emirates ^a
Bulgaria	Kiribati ^a	Nepal	Vanuatu ^a
Cabo Verde ^a	Korea, Dem. People's Rep. ^a	Oman ^a	Venezuela, RB
Cambodia ^a	Lao PDR	Panama	Vietnam
China	Lebanon	Paraguay	Yemen, (Arab) Rep. ^b
Comoros	Lesotho	Qatar ^a	Zambia
Czech Republic (Czechoslovakia) ^a	Liberia	San Marino ^a	
B) Closed Economies in 2001			
Afghanistan ^a	Korea, Dem. People's Rep. ^a	Nepal	Somalia ^a
Algeria	Marshall Islands ^a	Palau ^a	Timor-Leste
Bhutan	Micronesia, Fed. Sts.	Samoa ^a	Tonga ^a
Bosnia and Herzegovina ^a	Monaco ^a	San Marino ^a	Tuvalu ^a
Iran, Islamic Rep.	Montenegro ^a	São Tomé and Príncipe	Vanuatu ^a
Kiribati ^a	Nauru ^a	Serbia	
C) Closed Economies in 2015			
Afghanistan ^a	Marshall Islands ^a	São Tomé and Príncipe	
Iran, Islamic Rep.	Monaco ^a	Somalia ^a	
Korea, Dem. People's Rep. ^a	Palau ^a	Timor-Leste	

Notes: List includes all sovereign states that are members of the United Nations in 2019, if they existed at the time. For instance, in 1981, the USSR had not signed the GATT, but does not exist today so it is not on the list. Superscript (a) indicates that country is not included in analysis of sustained poverty reduction, either because it had kept the extreme poverty headcount at less than 3% of the population for the entire sample, or sufficient poverty data were not available. Superscript (b) indicates that economy was classified as open in 1981 by Sachs and Warner (1995) and Wacziarg and Welch (2008). Signatures are included only if in-force and have been notified to GATT or WTO.

