Trade Liberalisation and Regional Economic Performance in India: Coastal Gains, Interior Losses and the Anomalous Resilience of the Hills

Devashish Singh

Abstract

This paper examines how a common trade-liberalisation shock produces uneven regional outcomes across Indian states. A multi-region computable general equilibrium model, calibrated to a 2011 state-level social accounting matrix, is used to simulate a 20 per cent reduction in the export-tax wedge. Real household consumption rises in all three regional groups, but through different channels. Coastal states record the largest welfare and output gains, interior states experience a contraction in production that is more than offset by lower consumer prices, and the resilience of the hilly states is almost entirely price-led. A reconstruction of state-sector labour demand shows that national employment is fixed and that the shock primarily reallocates labour across sectors and destinations: coastal agricultural employment contracts while manufacturing and services expand, whereas the hilly increase is concentrated in agriculture. A manufacturing-employment-weighted cross-state regression further identifies a robust coastal-non-coastal divide. Adjusted manufacturing employment rises by 0.35 per cent in coastal states but falls by about 0.31 per cent in hilly states and 0.36 per cent in interior states. The regional contrasts remain significant after correction for multiple comparisons, while initial manufacturing and skill shares do not explain the difference. The results show that the regional incidence of liberalisation depends jointly on market access, production structure and household consumption patterns.

Keywords: trade liberalisation; regional inequality; computable general equilibrium; market access; labour reallocation; India.

1. Regional and Interstate Inequality: The Setting

Regional inequality in India is not a transient by-product of reform but a structural and durable feature of the national economy. Since the liberalisation of 1991, aggregate growth has accelerated sharply, with trend growth rising from the region of three to four percent in the pre-reform decades to eight or nine percent in the 2000s. Yet the gains from this acceleration have been distributed unevenly across space. Convergence across states, of the kind that neoclassical growth theory would predict, has largely failed to materialise, and interstate divergence rather than catch-up has been the dominant tendency of the post-reform period.

The most recent official accounting of this divergence is unambiguous. In their study of the relative economic performance of Indian states from 1960-61 to 2023-24, Sanyal and Arora (2024) show that the southern states, which were unexceptional performers before 1991, have emerged as the clear leaders since liberalisation. By 2023-24 the five southern states together accounted for close to thirty percent of national output, and the per capita income of every southern state had risen above the national average. The western littoral, led by Maharashtra and Gujarat, remains dominant, while much of the populous northern interior has slipped in relative terms. Liberalisation, in short, has coincided with a widening rather than a narrowing of the gap between the coastal and southern states on one side and the interior and eastern states on the other.

This geography of divergence has a recognisable shape. Coastal states with port access, deeper industrial bases and stronger integration into external markets have captured a disproportionate share of trade-related opportunity. Landlocked interior states, more dependent on agriculture and more distant from the gateways to world markets, have gained less and in several respects have fallen behind. The remote hill states of the Himalaya and the north-east occupy an ambiguous position: geographically peripheral and thinly integrated, yet not straightforwardly the largest losers from external opening. A substantial literature attributes these persistent gaps in part to unequal infrastructure endowments across states, with physical, social and financial infrastructure all found to track and reinforce interstate income differences (Timilsina, Sahoo, and Dash, 2022).

The specific puzzle that motivates this paper concerns the differentiated response of these three zones to a single national policy shock, namely a 20 per cent reduction in the export-tax wedge. The analysis points to three

patterns. Coastal port states gain the most, consistent with their superior market access and more diversified production base. Interior states experience relative or absolute losses, consistent with the specific-factors expectation that regions concentrated in import-competing activity, and constrained by limited factor mobility, bear the burden of adjustment. The remote hilly states, unexpectedly, prove resilient, and their resilience is led not by output but by prices: their comparatively industry-oriented consumption baskets mean that the fall in the prices of manufactured and traded goods raises real household purchasing power even where local production does not expand. The task of the analysis is to explain this pattern, and in particular to distinguish the output-led gains and losses of the coastal and interior states from the price-led gain of the hills.

The study pursues this question through a multi-region computable general equilibrium model of India, calibrated at the level of individual states and grouped, for interpretive purposes, into coastal port states, interior states and hilly states. The structure of that model, together with its calibration, closure and simulation design, is set out in full in a subsequent section. The present section has established the empirical setting. Section 2 reviews the literature within which the analysis is located, moving from the Indian evidence on openness and regional inequality, through the roles of factor immobility, market access and economic geography, to the treatment of the consumption and price channel, then to cross-country evidence from other large, regionally diverse economies, and closing with the methodological traditions in multi-region CGE modelling on which the study draws and a note on the scope for extension through geographic information systems.

2. Literature Review

The relationship between trade openness and inequality within a country is theoretically ambiguous. Standard trade theory offers competing predictions, and the empirical record is correspondingly mixed: openness has been found to widen regional disparities in some settings and to narrow them in others, with the outcome contingent on country characteristics, the mobility of factors and the initial spatial distribution of activity. The Indian literature, however, has increasingly concluded that in the Indian case specifically external opening has tended to aggravate rather than relieve interregional inequality. This review is organised around the mechanisms through which that outcome arises, since it is those mechanisms, and not the reduced-form association alone, that a general-equilibrium analysis of the regional question must represent.

2.1 Trade openness and regional inequality in India

A first strand documents the aggregate association between openness and interstate inequality. Daumal (2013), using time-series methods for India over 1980 to 2004, finds that India's trade openness significantly aggravates inequality among states, in contrast to her finding for Brazil, where openness is associated with reduced regional disparity. Barua and Chakraborty (2010) reach a compatible conclusion by locating the source of rising interregional inequality within the production structure of the economy, showing that interstate inequality is systematically related to cross-regional inequality in agriculture and manufacturing, and that trade openness is a key determinant of the manufacturing share of income across regions. The states already better placed in manufacturing capture the larger share of the gains from opening.

The wage evidence complicates any simple reading of these results. Mishra and Kumar (2005), evaluating the 1991 liberalisation through variation in industry wage premiums and trade policy across industries and over time, find a strong and robust negative relationship between tariff changes and wage premiums. Because tariff reductions were proportionately larger in sectors employing a greater share of unskilled workers, the pass-through of liberalisation-induced productivity gains to wages actually raised the relative incomes of unskilled workers, implying, on their reading, a reduction in wage inequality. The lesson for regional analysis is that the sign of a distributional effect depends on which margin is examined, the industry, the factor, the region or the household, and on whether the analysis captures general-equilibrium interactions across those margins. This is a central reason for approaching the regional question through a full general-equilibrium framework rather than a partial or reduced-form one.

A related literature stresses that a large part of Indian output and employment lies in the informal sector, so that the incidence of trade reform cannot be read off formal-sector variables alone. Sinha and Adam (2004), using a CGE model of India that distinguishes formal and informal activity, show that trade reforms transmit to households through informal labour and product markets in ways a formal-only model would miss. This work underlines the importance of resolving household income and consumption at a level of detail sufficient to capture heterogeneous transmission across regions.

2.2 Factor immobility and the incidence of adjustment

The most influential microeconometric evidence on the regional incidence of Indian trade liberalisation is that of Topalova (2007, 2010). Exploiting the 1991 reform and the variation across districts in pre-reform industrial composition, together with the variation across industries in the depth of liberalisation, she identifies the causal effect of tariff exposure on local outcomes through a difference-in-differences design. Rural districts more exposed to liberalisation, through the concentration of import-competing industries in their employment mix, experienced slower progress in poverty reduction, offsetting a material share of the poverty decline of the 1990s. The finding is consistent with a specific-factors model of trade and is attributed to the extremely limited mobility of labour and capital across regions and industries in India.

This mechanism is central to the regional question. If factors moved freely across states in response to adverse relative-price changes, the regional incidence of a national tariff cut would be arbitraged away and the effect on any given region would tend towards zero. It is precisely because labour and capital do not move freely that the shock leaves a durable spatial imprint, penalising regions locked into import-competing activity and rewarding those positioned in expanding traded sectors. Any credible model of the Indian case must therefore allow factor mobility across regions to be limited rather than frictionless, and must treat capital as substantially immobile in the short run. It is worth noting that the state-level evidence is not unanimous. Working at a more aggregated level and broadening the openness measure to include non-tariff barriers, other studies have reported more benign outcomes where labour markets are more flexible (Hasan, Mitra, Ranjan and Ahsan, 2012), which points to the degree of labour-market flexibility and the level of spatial aggregation as decisive parameters.

The Indian findings belong to a wider international literature that reads the incidence of trade shocks through the local labour market. In the most influential contribution, Autor, Dorn, and Hanson (2013) analyse the effect of rising Chinese import competition on local labour markets in the United States, exploiting variation across markets in initial industrial specialisation, and find that more exposed local markets experienced higher unemployment, lower labour-force participation and reduced wages, with the effect concentrated where import-competing manufacturing was located. The parallel with Topalova is close: in both cases a national trade shock produces sharply localised outcomes because workers and capital do not move freely enough to smooth the

incidence across space. This body of work established the local labour market, rather than the nation or the industry, as the natural unit at which to measure the distributional consequences of trade, and it is the analytical stance adopted here in taking the state as the unit of incidence.

The frictions that generate this localisation have since been measured directly. Artuc, Chaudhuri and McLaren (2010), estimating a rational-expectations model of intersectoral labour reallocation on United States data through an Euler-type equilibrium condition, recover switching costs that are strikingly large: in their preferred specification the mean cost of moving between broad sectors is roughly six and a half times average annual earnings, with a standard deviation of about five times, and no correction they apply brings the mean below four times annual earnings. The estimates carry a subtlety that reduced-form work cannot deliver. Because the idiosyncratic component of moving costs has high variance, gross worker flows exceed net flows by an order of magnitude, and this churning spreads the gains from liberalisation widely enough that, in lifetime welfare terms, an opening can leave even import-competing workers better off through the option value of mobility, even though the wage in the liberalised sector falls sharply at impact, overshooting its long-run decline, and never regains its former level. The result is not automatic: the option value fails to dominate at lower discount factors, and in their extension with worker heterogeneity it is age rather than education that decides who loses, with older import-competing workers the only group harmed. The welfare accounting of a trade shock must therefore follow workers through their reallocation choices rather than being read off the wage alone.

Three features of that framework bear directly on the present model. First, although their theory places no restriction on the distribution of the idiosyncratic benefit a worker derives from being in a given sector, beyond independence, identical distribution across workers, sectors and dates, and a mean of zero adopted as a normalisation, the functional form imposed in order to estimate the model is an extreme-value distribution, from which the gross flow between any two sectors follows as a multinomial logit share whose scale parameter is the dispersion of the shocks; the authors note that other distributions could in principle be used but would not in general deliver a closed form. The migration equation employed here has the same structure, allocating workers across sector and destination arcs in shares proportional to a constant-elasticity power function of the real return at the destination, which is that logit written multiplicatively. The difference lies in the index rather than in the

form of the choice, since theirs is a discounted continuation value net of a common moving cost while the version used here is a current real return deflated by state-specific consumption and housing prices and net of a mobility wedge. Their estimates therefore discipline the interpretation of the migration elasticity directly, the parameter governing the dispersion of their shocks occupying the position in the choice probability that the inverse of that elasticity occupies here: workers move a great deal, but the movements respond weakly to wage differentials, because idiosyncratic and nonpecuniary motives dominate the decision. Second, because the idiosyncratic benefits influence workers' choices, the long-run elasticity of intersectoral labour supply is finite, so that wages are not equalised across sectors even in the steady state, and a sector that loses protection must suffer a decline in its relative long-run real wage, since that is the only way it can sustain a net loss of workers. The counterpart in the present study is weaker and should be stated as such, for the model used here is comparative-static and its route-specific attractiveness parameters are calibrated to reproduce observed allocations at observed real returns, so that the initial differentials are lodged in those parameters rather than derived from the mobility friction. What survives the translation is the elasticity argument alone: mobility responds to real returns with a finite elasticity, so that liberalisation shifts the differentials rather than eliminating them, and the regional gaps reported below, which carry across from sectors to regions for the reason given next, are a property of imperfect mobility and not evidence of an adjustment left unfinished. Part of any such gap is in addition a compensating differential for housing rents, which have no counterpart in their framework, and their switching-cost estimates accordingly bear on the interpretation of these frictions rather than on their calibration, the cost they measure being paid once upon moving whereas the friction here is a standing wedge on the return to a route. Third, Artuc, Chaudhuri and McLaren note that the framework extends without amendment from industries to industry-region cells, and that they confine themselves to the sectoral application only because interregional mobility in the United States data is too thin to identify the parameters, a constraint that binds a fortiori in India, where interstate migration is lower still. The regional counterpart is supplied by Kennan and Walker (2011), who model interstate migration in the United States as an optimal sequence of location choices, estimated on National Longitudinal Survey of Youth panel data for white male high-school graduates among whom the annual interstate migration rate is 2.9 per cent. They find that income prospects influence migration substantially, operating both through geographic differences in mean wages and through a tendency to move in search of a better locational match when the

income realisation in the current location is unfavourable. Their estimate of the deterministic cost of moving is very large, about three hundred and twelve thousand dollars for the average mover in 2010 prices when the payoff shocks are set aside, but two qualifications must travel with that figure. The authors attribute its magnitude principally to the weakness of the observed relationship between income levels and migration probabilities rather than to the resource cost of relocating, and they show that once the shocks which actually trigger observed moves are taken into account, the average cost of the moves that are made is negative. More important for present purposes, they do not conclude that mobility is thereby rendered ineffective. Simulating a ten per cent change in the mean wage of a single state, they obtain spatial labour supply elasticities above one half for the large states they consider, with adjustment gradual but largely complete within a decade, and conclude that the supply of labour responds quite strongly to spatial wage differences despite the low observed rate of migration. What their evidence supports, then, is not the absence of interstate arbitrage but its operation through a finite elasticity and over a horizon of years, which is the property that a treatment of mobility as costly and imperfect is intended to capture.

The Brazilian liberalisation of the early 1990s has become the sharpest laboratory for the regional counterpart of these frictions. Kovak (2013) supplies the theoretical foundation for the region-level exposure measures used throughout this literature, deriving from a specific-factors model of regional economies the weighted average of liberalisation-induced price changes that constitutes a region's shock, with explicit guidance on the treatment of the non-traded sector, and finds that Brazilian regions facing a ten percentage point larger price decline experienced wage declines four percentage points larger. Dix-Carneiro and Kovak (2017) then trace the time profile: the effects of regional tariff exposure on formal earnings and employment did not fade but grew, with the impact after twenty years roughly three times the impact after ten, a pattern inconsistent with spatial-equilibrium models in which arbitrage erodes local shocks, and one they attribute to imperfect interregional labour mobility combined with labour-demand dynamics driven by slow capital adjustment and agglomeration economies. That combination, costly migration interacting with agglomeration externalities, is precisely the mechanism pair embedded in the present model, and their finding that regional divergence deepens rather than decays over two decades cautions against reading the comparative-static results here as a transitory phase that mobility will quickly undo.

The more benign state-level findings for India deserve elaboration, since they identify the conditioning role of institutions. Hasan, Mitra, Ranjan and Ahsan (2012), using state and industry-level data on unemployment and trade protection, find no evidence that the Indian reforms raised unemployment in the aggregate, and find that urban unemployment declined with liberalisation in states combining flexible labour-market institutions with larger employment shares in net-exporter industries. Two features of their study matter here. The heterogeneity of outcomes across states confirms that the incidence of a common national shock is mediated by state-level structure, which is the premise of any multi-region treatment; and their explicit defence of the state as an independent labour market, grounded in the considerable regulatory powers that India's constitutional arrangements assign to the states and in the low rates of interstate migration documented in the demographic literature, is the same premise on which the regional units of the present model rest. The mechanism behind the institutional conditioning has a structural counterpart: Kambourov (2009) shows, in a calibrated model in which sectoral reallocation is costly because of sector-specific human capital and firing costs, that firing costs exert a large negative effect on the gains from trade reform, which is precisely the margin along which Indian states differ.

Between the regional studies and the individual household sits a literature that measures the incidence of trade reform at the household level by passing general-equilibrium price and wage changes through survey data, and its findings anticipate the consumption-price channel taken up in Section 2.5. Winters, McCulloch and McKay (2004), surveying the evidence, conclude that no simple general statement links liberalisation to poverty: the incidence runs through consumer and producer prices, through wages and employment, and through government revenue and spending, channels that can pull in opposite directions for the same household, even though the balance of evidence suggests liberalisation tends to relieve poverty in the long run. Attanasio, Goldberg and Pavcnik (2004) document the earnings side for Colombia, where industry wage premia fell most in the sectors facing the deepest tariff cuts, sectors that were initially low-wage and intensive in unskilled labour, so that the pressure of reform bore hardest on workers attached to the liberalised industries, although they judge the overall effect of the reforms on the wage distribution to have been small. Chen and Ravallion (2004) make the mechanics concrete for China's WTO accession, mapping general-equilibrium price and wage changes onto national household surveys, and find impacts on aggregate poverty and inequality that are negligible even as the

incidence across household types and regions is sharply heterogeneous, rooted in differences in consumption behaviour and income sources, with direct implications for compensatory policy. Porto (2006) develops the corresponding method for Argentina and finds that the tariff reforms of Mercosur benefited poor and middle-income households, with gains of around six per cent of initial expenditure at the bottom of the distribution. For India, Ural Marchand (2012) estimates the pass-through of tariffs to domestic prices separately for rural and urban areas and lets wages adjust in equilibrium, finding welfare gains across the entire per capita expenditure distribution, generally pro-poor and driven substantially by the consumption of tradables. The consistent message of this literature, that the same reform is experienced differently by households according to what they earn and what they buy, is the household-level statement of the decomposition this paper performs at the level of states.

2.3 Economic geography, market access and the coastal advantage

Why coastal and port-proximate regions gain disproportionately is the domain of the new economic geography and its quantitative descendants. The organising concept is market access, which in this literature carries a precise meaning. Its ancestor is the market-potential function of Harris (1954), a descriptive construction in which the demand available to any location is the distance-weighted sum of purchasing power everywhere else. Fujita, Krugman and Venables (1999) supplied the theoretical foundation, and in the form employed by Redding and Venables (2004) the market access of a region is the sum across all destinations of their expenditure, each destination weighted by its own price index and by the cost of shipping to it, this composite entering the equilibrium condition that fixes the highest wage a firm in that region can afford to pay. A companion term, supplier access, aggregates the supply capacities of all origins weighted by the cost of shipping from them, and governs the local price index rather than local demand. The distinction is worth preserving, since the two operate through different channels, one upon the revenue a region’s producers can earn and the other upon the cost of what it buys, and a region well placed on both counts faces a lower effective trade cost on either side of its accounts. Opening therefore raises the derived demand for the goods and factors of well-connected regions by more than for remote ones. Redding and Venables (2004) estimate a structural economic-geography model on cross-country data and show that access to markets and sources of supply is a statistically significant

and quantitatively large determinant of per capita income differences, operating through exactly the mechanisms the theory emphasises. Hanson (2005) provides complementary micro-evidence, estimating an augmented market-potential function derived from the Krugman model in which wages rise with proximity to consumer markets, and recovering parameter values that confirm the joint importance of scale economies and transport costs, though with demand linkages that are strong but geographically bounded.

Two aspects of that framework bear emphasis here. The first is why differences in access should translate into differences in the location of activity at all, which is the question Krugman (1991) answered by showing that the interaction of increasing returns at the level of the plant, transport costs and factor mobility can generate a core-periphery structure endogenously, manufacturing concentrating in the region that begins with the larger market and the initial advantage compounding through the demand and cost linkages that follow. Access on this account is not a passive description of geography but a force that reproduces itself, which is what distinguishes the economic geography tradition from a simple transport-cost story. The second is that market access divides into a domestic and a foreign component, since a region draws demand both from the remainder of its own country and from abroad, and Redding and Venables (2004) find that the domestic component accounts for a substantial part of the cross-country variation in income per head. For a model of Indian states, whose traded output is sold partly to other states and partly overseas, that decomposition is not incidental. A state may enjoy strong access to the domestic market while remaining remote from any port, and the incidence of an external shock upon it will differ accordingly from the incidence upon a state well placed on both margins.

The mechanism has been formalised most directly for the coastal question by Cosar and Fajgelbaum (2016), who introduce an internal geography into the canonical comparative-advantage model of trade to study the regional effects of external opening. Their model generates a dual-economy structure in which locations near international gateways specialise in export-oriented sectors while more distant interior locations trade little with the rest of the world, and it rationalises the observed patterns of specialisation, employment and relative income in developing countries that have opened to trade. This is, in essence, the coastal-interior divide at the heart of the present study rendered as theory: proximity to the gateway determines whether a region is drawn into external

trade or left oriented towards the domestic market, and it is this differential exposure, rather than any difference in policy, that produces divergent regional fortunes under a common liberalisation.

For India specifically, the magnitude of these internal frictions has been quantified by Van Leemput (2021), who builds a multi-sector trade model with nonhomothetic preferences in which Indian states trade both with one another and internationally, and disciplines it with state-level price-dispersion and trade-flow data. Two of her findings bear directly on the present analysis. First, internal trade barriers in India are large, amounting on average to about forty percent of total trade cost, and they vary systematically across states with distance to the nearest port, which is precisely the structural source of the coastal advantage modelled here. Second, the welfare gains from domestic integration exceed those from eliminating external import barriers, a result that underlines how far the regional incidence of trade policy is mediated by internal geography. Together these studies establish that in India the effective trade cost faced by a state, and hence its exposure to any external shock, is a function of its position in the internal transport geography.

Two features of this literature bear on the design of a regional model. The first is the coupling of market access with increasing returns to scale: it is the interaction of scale economies and trade costs, not either alone, that generates the spatial concentration of activity and the disproportionate gains of well-placed regions. This suggests that increasing returns should operate in the sectors where scale is empirically salient, manufacturing and services, while agriculture is treated as closer to constant returns. The second is the theoretical foundation of the gravity structure of trade, formalised in the Ricardian model of Eaton and Kortum (2002), which yields tractable structural equations for bilateral trade in terms of technology, comparative advantage and geographic barriers. Both point towards a trade structure that ties interstate and international flows to relative prices and trade costs in a manner consistent with each region's geographic position.

2.4 Quantitative spatial models and labour-market dynamics

The recent quantitative spatial economics literature has integrated goods trade, costly labour mobility and geography into unified general-equilibrium frameworks, and it is this tradition that most directly informs the treatment of the labour market in a regional trade model. Redding (2016) develops a quantitative spatial model combining a rich geography of trade costs with labour mobility and heterogeneous worker preferences across

locations, and shows, importantly for the present question, that factor mobility introduces quantitatively relevant differences between the counterfactual predictions of constant-returns and increasing-returns models. The welfare gains from trade in such settings depend not only on changes in domestic trade shares but also on the reallocation of population across locations. Redding (2020) surveys this body of work, distinguishing first-nature geography, the natural endowments of a location, from second-nature geography, the location choices of agents relative to one another, and documents how quantitative spatial models rationalise the reduced-form findings on the distributional consequences of trade.

The foundations of this class of models were laid by Allen and Arkolakis (2014), who develop a general-equilibrium framework capable of determining the spatial distribution of economic activity on a surface with essentially arbitrary geography. Combining the gravity structure of trade with labour mobility, they establish conditions for the existence, uniqueness and stability of a spatial equilibrium and derive a compact set of equations linking activity in each location to the geography of the surface, which they then use to estimate the topography of trade costs, productivities and amenities across the United States. Two features of their analysis carry directly into the design of the present model. The first is that the fortunes of a location are determined by its position within the entire configuration of the network and not by its own endowments alone, so that the exposure of a state to a national trade shock is a property of its relation to ports, to markets and to its neighbours. The second concerns the treatment of increasing returns. Allen and Arkolakis show that a unique and stable spatial equilibrium requires the strength of the agglomeration forces to be bounded relative to the dispersion forces generated by trade costs and congestion, since a sufficiently powerful externality admits multiple equilibria and leaves counterfactual exercises indeterminate. This result is the justification for the modest Marshallian externality adopted here, which is confined to industry and services and calibrated at a magnitude small enough to preserve a well-behaved equilibrium while still allowing scale to reward the states that expand. Redding and Rossi-Hansberg (2017) survey the literature that has grown from these foundations and set out the menu of building blocks from which such models are assembled, among them the treatment of many heterogeneous locations, gravity relationships governing both goods trade and the reallocation of labour, and exogenous or endogenous amenities, and they demonstrate the apparatus on policy exercises that reduce external and internal trade costs in turn, which is the precise pairing at issue in the Indian case.

How such a structure transmits a localised shock is shown by Caliendo, Parro, Rossi-Hansberg, and Sarte (2018), who trace the propagation of disaggregated productivity changes through intersectoral and interregional trade linkages in the United States and recover the aggregate, regional and sectoral elasticities of measured total factor productivity, output and employment to shocks originating in particular regions and industries. Their central finding is that these elasticities vary substantially according to which sector and which region is affected, and that the variation is governed by the spatial structure of the economy rather than by the size of the initiating disturbance. The corollary for the present study is that the regional incidence of a national trade shock cannot be read off the exposure of each state taken in isolation, since the outcome for any one state depends on the sectoral composition of its trading partners and on the linkages through which the shock reaches it. Their counterfactual exercises, which include experiments in transport infrastructure alongside productivity shocks, also establish the practice of using a calibrated regional model to conduct mechanism-specific experiments, an approach followed in the analysis below.

The dynamic counterpart most relevant to a trade shock with spatially uneven incidence is Caliendo, Dvorkin, and Parro (2019), who build a dynamic trade model with spatially distinct labour markets facing differing exposure to international trade, incorporating labour-mobility frictions, goods-mobility frictions, geography and input-output linkages. Calibrated to sectors and states, their analysis of the China trade shock shows that although the economy gains in the aggregate, trade and migration frictions cause welfare and employment effects to vary widely across local labour markets. This is precisely the pattern a study of Indian tariff reduction must be able to reproduce: aggregate gains coexisting with sharply differentiated regional outcomes, the differentiation arising from frictions in the movement of goods and workers. Taken together, this literature motivates a treatment of interstate labour reallocation as costly and imperfect, governed by the relative real returns to work in each destination net of mobility barriers, rather than as either frictionless or absent.

2.5 The consumption-price channel and the hill-state anomaly

The final substantive strand concerns the consumption side of trade adjustment, and it speaks most directly to the anomalous resilience of the hill states. Much of the regional-incidence literature is framed in terms of income, output and factor earnings, that is, the production side. Yet trade liberalisation affects households

through two distinct routes: as earners, whose factor incomes rise or fall with the fortunes of local industry, and as consumers, whose real purchasing power depends on the prices of the goods they buy. Where a region's consumption basket is weighted towards traded manufactures whose prices fall when tariffs are cut, households can gain in real terms even if local production stagnates. Topalova (2010) extends the district analysis to consumption across the income distribution, and the broader trade-and-poverty literature has increasingly stressed that the cost-of-living channel can offset, or even reverse, the sign of the production-side effect.

The theoretical foundation for this channel is set out by Fajgelbaum and Khandelwal (2016), who develop a method for measuring the unequal gains from trade that arise on the consumption side, on the reasoning that individuals consuming different baskets of goods are differentially affected by the relative-price changes that trade brings about. Their central finding, that trade tends to favour those who concentrate their spending in the more traded sectors, is the general principle of which the hill-state result is a regional instance. The distinction they draw, between the earnings channel that has dominated the trade-and-inequality literature since Stolper and Samuelson and the expenditure or cost-of-living channel that their method isolates, is exactly the distinction between the output-led and price-led sources of regional welfare change that the present study seeks to separate. Their framework supplies the analytical justification for resolving welfare at the level of the state-specific consumption basket rather than at the level of factor income alone.

The channel has a direct Indian general-equilibrium precedent. Pradhan and Sahoo (2008), in a twenty-eight sector CGE model of India with nine occupational household groups and a poverty line endogenised to consumer prices, find that a complete removal of import tariffs lowers factor incomes, household disposable incomes and welfare, yet reduces measured poverty in both rural and urban areas, because the across-the-board fall in consumer prices lowers the cost of the poverty basket by enough to lift households near the line out of poverty even as incomes decline. Their summary of the mechanism is exactly the decomposition pursued here: disposable incomes and consumer prices fall together, the decline in welfare is governed by the income side, and the decline in poverty by the price side. Two further findings are instructive. Relaxing import quotas on agricultural and food products, a shock transmitted to households principally through the consumption basket rather than the earnings base, raises welfare for almost all household groups and reduces poverty most in rural areas; and when the tariff

removal is instead made revenue-neutral through a compensating rise in domestic indirect taxes, the consumption-price gain is eroded, welfare falls for every group, poverty rises, and rural households, whose dependence on agriculture parallels that of the interior states in the present regional setting, bear the larger burden. The study establishes, for India and within a general-equilibrium accounting, that the same trade shock can carry opposite signs on the income and expenditure margins, and that the fiscal instrument chosen to accompany liberalisation can reverse the distributional outcome.

This is the mechanism that accounts for the hill-state result. The remote hilly states are not major producers of traded manufactures and gain little on the output side from a tariff cut. Their consumption baskets, however, are comparatively industry-oriented, as the underlying expenditure data confirm, and they are substantially reliant on goods brought in from outside the state. When the tariff reduction lowers the prices of manufactured and imported goods, the state-level cost of living in the hill states falls by more than the fall in their nominal incomes, so that real household consumption improves. The resilience of the hills is therefore price-led rather than output-led. Capturing it requires a framework that represents the consumption basket and the composite price of goods at the level of the individual state, so that the change in real consumption can be traced to its underlying commodity-price and basket-composition components. An analysis confined to the production side would misclassify these states as neutral or adverse cases and would miss the channel entirely. Isolating this price-led welfare gain, and distinguishing it cleanly from the output-led gains of the coastal states and the output-led losses of the interior, is the central analytical task of the study.

A challenge to this reading must be confronted directly, and it comes from the most careful available evidence on how international price changes reach consumers inside a developing country. Atkin and Donaldson (2015) assemble retail price data for narrowly defined, barcode-equivalent goods across 103 market towns in Ethiopia, 36 state capitals in Nigeria and 1,881 counties in the United States, together with newly collected information, obtained by interviewing the producing firms themselves, on the factory or the port of entry from which each product in the sample originates. Restricting attention to genuine trading pairs and using estimated pass-through rates to purge spatial price gaps of the mark-ups charged by intermediaries, they find that the effect of log distance on trade costs within their two African countries is between three and a half and five and a quarter

times its value in the United States, a ratio that narrows only to between two and a half and four when the distance metric is replaced by quickest-route travel time, which adjusts for the availability and the quality of roads together. In ad valorem terms, the additional cost of serving a location five hundred miles from the source rather than fifty miles from it is about twenty per cent of the mean product price in Ethiopia and twelve per cent in Nigeria, against four per cent in the United States. Pass-through from origin to destination prices averages approximately one half and declines significantly with distance from the origin in all three countries, so that retail prices in remote locations respond only weakly to a cost shock at the port or the factory gate. Because the pass-through rate also governs the division of surplus between consumers and intermediaries, they conclude that intermediaries capture the majority of the gain released by a fall in world prices and that their share rises with remoteness. A final result operates on the extensive margin: in the more distant locations the good is frequently not stocked at all, the probability of a product being found by price enumerators falling by twenty-two per cent in Ethiopia and eleven per cent in Nigeria between the most proximate and the most remote markets in the sample.

Set against the framework used here, the challenge is specific rather than wholesale, and stating it precisely matters. The present model is not spatially uniform in its treatment of traded-goods prices. Each state faces an ad valorem cost wedge on its international transactions, applied symmetrically to the price at which imports arrive and the price at which exports are cleared, and it is this wedge, together with state-specific Armington weights on imported and domestically sourced varieties, that generates the coastal advantage the model reports and that makes the composite consumption price differ across states in the first place. What the model does not represent is the interstate leg. Goods produced within the country enter every state at a pooled national price, so that a state buying from another state bears no cost for the distance between them and a change in that price arrives everywhere in full. The objection therefore reduces to two identifiable omissions rather than to a general charge of spatial uniformity: the absence of a domestic interstate distance wedge, and the imposition of complete pass-through in a setting where the measured rate is closer to one half and falls with remoteness. The first of these bears with particular force on the hill states, whose industry-oriented baskets are supplied predominantly by producers in other states rather than directly from a port, so that the leg on which their price-led gain depends is precisely the leg the model leaves frictionless.

Three considerations bound the force of this objection without dismissing it. The first concerns which of their three countries is the relevant comparator, since the quantities that matter here differ sharply across the three. Ethiopia and Nigeria sit at the severe end of the distribution of internal trade frictions and the United States at the other. The coefficient of the adjusted price gap on log distance is 0.011 in the United States against 0.037 in Ethiopia and 0.056 in Nigeria; and, more directly to the point, the share of total surplus accruing to consumers rather than to intermediaries declines with distance by a statistically significant margin in both African countries but not in the United States, where the estimated gradient is small and insignificant and implies a difference of roughly one percentage point between the least and the most remote locations, against four points in Ethiopia and thirteen in Nigeria. India lies between the two poles. Van Leemput (2021) places Indian internal trade barriers at about forty per cent of total trade cost, substantial but well short of the African benchmark, and Ural Marchand (2012), estimating tariff pass-through into Indian domestic prices separately for rural and urban areas, still recovers welfare gains across the whole per capita expenditure distribution driven substantially by the consumption of tradables. Direct evidence on the interstate leg itself is supplied by Donaldson (2018), who infers trade costs from archival data on the colonial railway network and finds that the railways reduced trade costs and interregional price gaps, raised interregional and international trade, and increased real income, districts gaining access recording a rise in real income of about sixteen per cent while districts bypassed by the network suffered a negative spillover. That evidence establishes both that an interstate wedge exists in India and that it responds to infrastructure, which places the mechanism the present model omits within Indian data rather than African. The quantity at issue for India is accordingly a few percentage points of surplus share rather than the reversal of a sign.

The second consideration concerns what their price sample measures and in which direction it travels. Their incidence exercise is a fall in the port price of an imported product, transmitted inland from the point of entry, whereas the shock studied here works outward, from producers in the interior towards the port, and attenuation measured on the inbound leg need not be symmetric on the outbound one. Their estimates are in any case drawn overwhelmingly from domestically produced goods, only five of fifteen Ethiopian and three of eighteen Nigerian products in the sample being imported. The sample is also, as the authors note, restricted to branded products carrying detailed descriptions and is on that account unlikely to be representative of consumer spending; it consists of packaged manufactures with high value-to-weight ratios and omits the staples, services and housing

that carry substantial weight in Indian state consumption baskets. The gradient they estimate applies to a subset of the industrial composite rather than to the basket as a whole. The third consideration is that the objection speaks to the magnitude of the price-led term and not to its sign, nor to the validity of the decomposition. Estimated pass-through rates are almost uniformly positive, so a fall in the composite price of traded goods still lowers the state-level cost of living, and the composition of the state consumption basket, which is the second factor in the price-led term and the one that distinguishes the hill states from the interior, is untouched by the rate of pass-through. It is likewise the gradient rather than the level of their incidence result that should be carried across. The division of surplus between consumers and intermediaries is recovered under an assumption that the authors themselves describe as particularly stark, namely that demand curvature is constant across locations while the competitiveness of the trading sector is constant across products, and it rests on a normalisation that treats the least competitive location in each country as a monopoly. The underlying pass-through estimates are noisy, are winsorised from below, and are conceded to be biased upward wherever demand shocks are spatially correlated. What survives robustly is the direction of the gradient and its relative steepness across countries, which is what the present argument requires.

The appropriate response is to treat the price-led gain reported for the hill states as an upper bound and to convert the missing assumption into a parameter that can be varied. Replacing the pooled national price of domestically sourced goods with a state-specific delivered price, reached through a distribution margin rising in the log of distance from the supplying region, would do this directly, and Atkin and Donaldson supply both the functional form for such a margin and a bracketed range for its steepness, their United States and African estimates plausibly bounding the Indian case from below and from above. This is the endogenous, distance-based trade cost set out among the extensions in Section 2.8, and it identifies the sensitivity analysis that would settle the matter, since the hill-state result survives if and only if the price-led gain net of a plausible interstate wedge still exceeds the output-led loss. A further implication of their work lies outside the reach of the present framework altogether. If a substantial share of the surplus released by a trade shock is captured by intermediaries rather than by households, then a model with competitive distribution and no trading sector overstates the consumer incidence of the shock in every state, and most of all in the remote ones. That is a limitation of the

competitive Armington representation rather than of the decomposition, and it is recorded here as a boundary on the interpretation of the welfare results that follow.

2.6 Cross-country evidence from other large, regionally diverse economies

The Indian pattern is not unique. The same broad dynamic, in which coastal or gateway regions capture the gains from opening while interior regions bear the adjustment, recurs across large and regionally diverse economies, and multi-region studies in those settings corroborate both the pattern and the mechanisms emphasised above. Reviewing them serves two purposes: it establishes that the Indian result is an instance of a general regularity rather than a national peculiarity, and it identifies the specific modelling choices that other authors have found necessary to reproduce it.

The general regularity is established most directly by the cross-country econometric evidence. Rodriguez-Pose and Gill (2006), comparing the evolution of trade and regional disparities across eight major economies, among them India, find that internal disparities are considerably larger in the developing countries of their sample than in the developed ones, and that the link between trade and regional inequality emerges most clearly once the sectoral composition of trade is taken into account: as primary products lose weight in the composition of total trade, regional inequalities tend to rise, a shift that concentrates the gains from an increasingly industrial trade in the regions that host the expanding traded sectors, and one whose burden falls disproportionately on developing countries, where internal disparities are larger to begin with and where the agricultural share of trade has been declining fastest. This compositional mechanism is the cross-country counterpart of the finding of Barua and Chakraborty (2010) for the Indian states, and it bears directly on a simulation in which traded manufacturing is the locus of both the increasing returns and the coastal advantage.

Rodriguez-Pose (2010) sharpens the conditionality in a panel of twenty-eight countries over 1975 to 2005. Increases in trade per se do not raise territorial polarisation; trade interacts instead with country-specific conditions, and the greatest rises in territorial inequality are observed where interregional differences in sectoral endowments are large, where the share of government expenditure is low, and where high internal transaction costs coincide with a close correspondence between the regional distribution of income and regional access to foreign markets, a configuration more characteristic of low and middle income countries. Each of these

conditions describes the Indian case as represented here: the sectoral structures of the coastal, interior and hilly zones differ sharply; internal trade costs are of the order documented by Van Leemput (2021); and the states with the best access to external markets are, for the most part, also the richest. The cross-country evidence therefore places India squarely in the class of countries in which external opening should be expected to widen regional gaps. The conditioning role of government expenditure carries a further implication for a federal polity: Rodriguez-Pose and Gill (2004) read the worldwide movement towards devolution as a weakening of the equalisation function of national government, with the costs of adjustment falling most heavily on the regions least equipped to bear them, a consideration that bears on whether the interior losses identified in the simulation would be moderated or entrenched by the prevailing fiscal architecture.

China is the paradigmatic case of post-reform coastal-interior divergence. Wan, Lu, and Chen (2007) show empirically that globalisation has been a significant force widening regional income inequality within China, with the coastal provinces pulling away from the interior over the reform period. A recent quantitative study in Economic Modelling (2026) sharpens the mechanism: building a model with tariff barriers and labour migration, it finds that coastal zones disproportionately capture export-driven growth through lower transport costs and agglomeration economies, while inland regions face intensified import competition and stagnant labour demand behind high domestic trade barriers, and that internal labour mobility moderates the whole distribution by drawing workers from the interior towards the export-oriented coast. That is the Indian result reproduced in another setting, and it independently supports the two mechanisms on which the present model rests, namely the port-access wedge that advantages coastal states and the costly, incomplete reallocation of labour across regions. Its emphasis that spatially embedded trade costs, rather than tradable-factor endowments, drive the divergence is the same conclusion the coastal and interior results here point to.

Mexico provides the closest architectural precedent for a multi-region model with migration. Hinojosa-Ojeda, Robinson, and De Paolis (1999) construct the GNAFTA-CGE model, in which sub-regional CGE models of Mexico, Central America, the Caribbean and the United States are linked through trade and migration flows, with several labour categories including rural labour, urban unskilled, skilled and white-collar workers. Their framework is the direct antecedent of the structure adopted here, in which regions are distinct economies joined

by trade and by a costly reallocation of labour, and the Mexican experience under North American integration furnishes the classic north-south analogue of a gateway-versus-interior divide widening under liberalisation.

The archipelagic case of Indonesia sharpens the port and transport dimension that is central to the coastal advantage. Using the interregional IndoTERM model, Horridge, Yusuf, Ginting, and Aji (2016) simulate reductions in road and sea transport costs and find that peripheral regions are poorer precisely because they face higher transport costs, so that the core regions capture the bulk of the gains while measurable spillovers reach the periphery. Because IndoTERM represents sea transport and port efficiency explicitly, it is the natural precedent both for the port-access wedge in the present model and for the geographic-information-system extension discussed below, and it belongs to the same bottom-up TERM lineage on which the modelling approach here draws.

The long run of Spanish economic history supplies a historical depth that the contemporary cases lack. Paluzie, Pons and Tirado (2004), constructing indices of the geographical concentration of manufacturing across Spanish regions at two levels of aggregation over the period 1856 to 1995, confirm that the integration of the national market stimulated the spatial concentration of industry, which gathered above all in the coastal regions of Catalonia and the Basque Country, and they stress that the relationship between integration and concentration is neither unique nor stable, and may trace an inverted U over the very long run, with concentration rising through the first century of integration before moderating in the closing decades of the twentieth. Two lessons follow for the present exercise. The first is that the association between falling trade costs and the spatial concentration of industrial activity, on which the increasing-returns mechanism of the model rests, is not an artefact of recent globalisation but a durable regularity visible over a century and a half. The second is a caution: the Spanish experience indicates that the concentrating phase need not be permanent, so that the regional divergence generated by liberalisation in the simulation is better read as the impact phase of an adjustment whose eventual moderation depends on forces, among them congestion, factor mobility and policy, that lie partly outside a comparative-static model.

Taken together, the cross-country panel evidence and the Chinese, Mexican, Indonesian and Spanish cases yield a consistent lesson. Across very different institutional settings, it is internal geography and trade cost, rather

than factor endowments alone, that determine which regions gain from external opening, and it is the limited mobility of factors that converts a common national shock into a durable spatial pattern. This is precisely the regularity the present model is built to reproduce for India, with the additional contribution, not attempted in these studies, of separating the price-led source of a region's welfare change from the output-led source.

2.7 Methodological traditions in multi-region CGE modelling

A study of this kind sits within the tradition of multi-region computable general equilibrium analysis, and it is useful to set out the methodological literature on which it draws before the model itself is described. The single most direct antecedent is Naranpanawa and Arora (2014), who construct what they describe as the first single-country multi-region CGE model of the Indian economy specifically to examine the trade-liberalisation and regional-disparity linkage. Their central result, that in the short run liberalisation benefits the richer, faster-growing middle-income states and has a marginal or negative effect on the poorer states, is the natural point of comparison for the present study, and their motivation, that partial-equilibrium and reduced-form approaches cannot capture the economy-wide, multi-channel transmission of a trade shock, is also the methodological justification for a general-equilibrium approach.

The spatial CGE programme associated with Haddad and Hewings and their collaborators, developed principally for Brazil and later extended to other economies, supplies much of the apparatus for handling geography within such a model. This body of work is important on three counts. First, it demonstrates how to embed non-constant returns to scale and non-iceberg transport costs within an operational interregional CGE model, showing that in a spatially concentrated economy the assumption of constant returns is untenable and that scale economies allow dominant regions to exploit their position asymmetrically (Haddad and Hewings, 2005). Second, it establishes the interregional CGE method for tracing how transport costs, market access and scale interact to shape regional growth and inequality (Haddad and Hewings, 1998, 2004). Third, and most directly relevant to the coastal dimension of the regional question, Haddad, Hewings, Perobelli, and dos Santos (2010) model port infrastructure explicitly as a barrier to trade, showing that the efficiency of ports of entry and exit propagates through the hinterland and materially conditions the regional distribution of the gains from openness. The extension to Egypt integrates the interregional CGE model with a geo-coded transport network to assess

domestic integration (Haddad, Lahr, Elshahawany, and Vassallo, 2016), and the earlier Brazilian studies apply the same apparatus to the regional effects of trade agreements (Haddad, Domingues, and Perobelli, 2002).

The construction of the multi-region database itself, in which each state is represented as a distinct economy with its own production, consumption and trade structure and is then linked to the others through interstate trade, follows the bottom-up strategy exemplified by the TERM model of Horridge (2011) and, in an earlier and more aggregated form, by the Canadian regional model of Jones and Whalley (1989). This bottom-up approach, as distinct from a top-down regionalisation of national results, is what permits genuinely state-specific outcomes to emerge from state-specific structures, which is essential when the object of interest is precisely the divergence of regional responses. On the representation of the labour market, Boeters and Savard (2013), in their survey of labour-market modelling in CGE frameworks, catalogue the available options on the supply side, the demand side and in wage formation, and make clear that the treatment of labour is rarely neutral for distributional results. The wider methodology of CGE construction, calibration and closure follows the conventions codified in the Handbook of Computable General Equilibrium Modeling (Dixon and Jorgenson, 2013).

Finally, the concern with the labour-market and welfare incidence of trade liberalisation places the study alongside a set of applied CGE analyses of developing economies. Galrao Carneiro and Arbache (2003) use a CGE model to assess the labour-market impact of trade liberalisation in Brazil, framing the question through the Heckscher-Ohlin and Stolper-Samuelson mechanisms and finding net welfare gains alongside heterogeneous sectoral effects, while the multi-country CGE studies of trade liberalisation and poverty assembled for African and Asian economies (Bussolo and Round, 2006; Cockburn, Decaluwe and Robichaud, 2008) apply the same apparatus to household welfare under reform. These traditions provide the methodological foundation for the approach adopted here; its specific structure, which adds the explicit interstate spatial dimension, the costly interstate reallocation of labour and the state-level consumption-price decomposition that the Indian regional question requires, is described in a subsequent section.

3. The Interregional CGE Model

3.1 Overview

The analysis rests on a comparative-static interregional computable general equilibrium model of the Indian economy, calibrated to a benchmark year of 2011. The model recognises twenty-four states, three production sectors covering agriculture, industry and services, and two categories of labour distinguished by skill. Results are generated bottom up, in the sense that national outcomes are the aggregate of the individual state outcomes rather than the product of a national model to which regional shares are subsequently applied. Agents optimise at the state level, so that the structure of each state economy, its sectoral composition, its factor endowments and the composition of the basket its households consume, is allowed to condition its response to a shock that is common to the country as a whole.

The architecture follows the interregional tradition represented by Haddad and Hewings (2005) and, most closely in its treatment of accessibility, by Haddad, Lahr, Elshahawany and Vassallo (2016), whose presentation of the nesting structure is adopted below. Three departures from that template should be recorded at the outset, since each is motivated by the question this study addresses. First, the domestic commodity is traded on a single national market rather than through bilateral flows between every pair of states, so that the sourcing nest which those models place at the foot of the production and consumption trees is absent here, and spatial differentiation enters instead through the international leg. Second, labour is reallocated across states as well as across sectors, through a discrete-choice share equation in the tradition discussed in Section 2.2, rather than through the wage-relativity closures customary in the Australian tradition. Third, housing is in fixed supply in each state and enters both the cost of living and the return that governs the migration decision, so that the congestion cost of in-migration is borne explicitly. A fourth feature, external economies of scale in industry and services, is discussed with the production structure below.

The model is written in levels and solved as a square system of simultaneous non-linear equations rather than as an optimisation problem or as a linearised system in percentage changes. This differs from the Johansen-class implementation used by Haddad and his co-authors, and the practical consequence is that the results

reported here are exact for the size of shock administered rather than a first-order approximation to it, which matters where the shock is large and where increasing returns make the response non-linear.

3.2 Production structure

Figure 1 sets out the nesting structure of production. Dotted boxes identify the functional form governing substitution at each level of the hierarchy, and producers in each state and sector choose input requirements per unit of output by minimising cost subject to that technology.

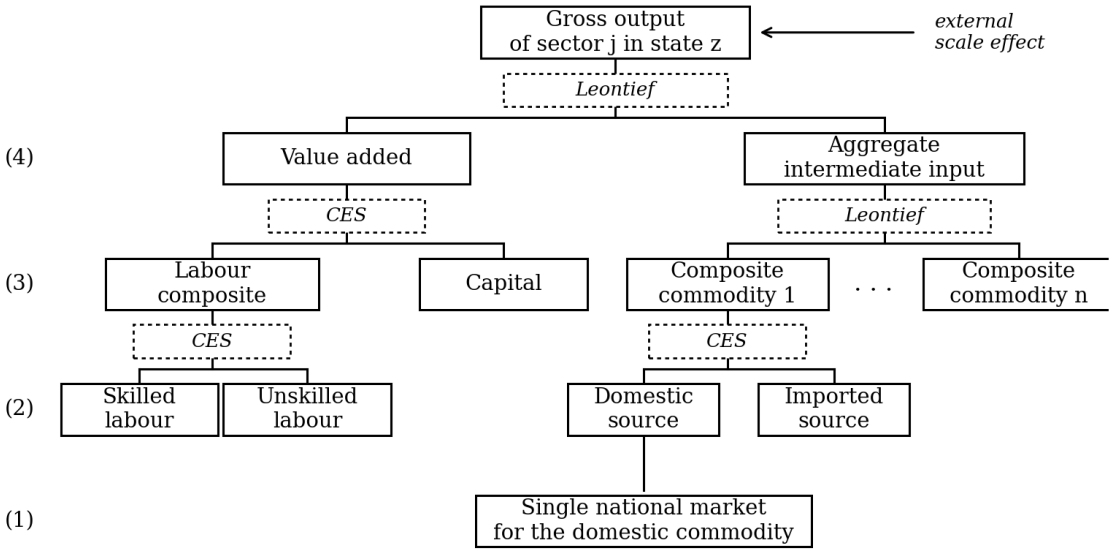


Figure 1 Nesting structure of state production technology

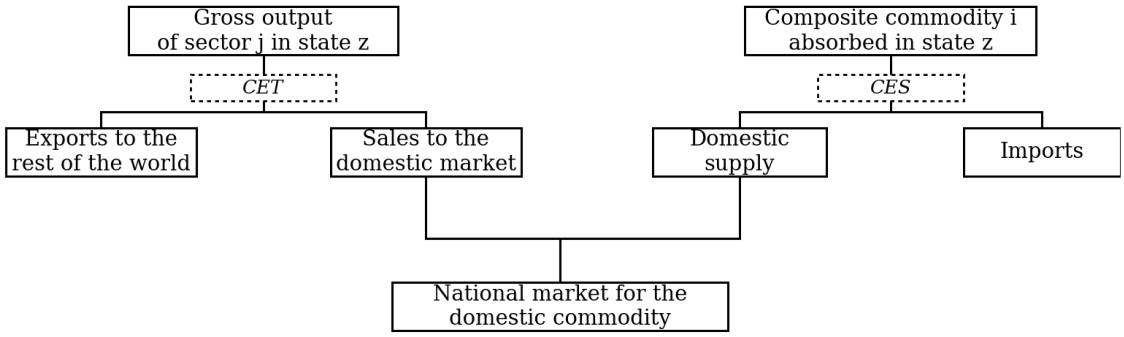
At the top level, gross output is produced from value added and an aggregate intermediate input in fixed proportions, so that neither can be substituted for the other and the intermediate input ratio is a technological constant of the state and sector. At the second level the aggregate intermediate input is itself a Leontief composite of the individual commodities, with input-output coefficients that vary across states as well as across using sectors, while value added is a constant-elasticity-of-substitution composite of a labour aggregate and capital. Cost minimisation at this level yields a demand for the labour aggregate relative to capital that depends upon the ratio of the capital rental rate to the sectoral wage composite, raised to the elasticity of substitution between the two factors. At the third level the labour aggregate is a further constant-elasticity composite of skilled and

unskilled labour, so that the relative demand for the two skills responds to the ratio of the composite wage to the wage paid to each type, and each individual commodity entering the intermediate composite is an Armington aggregate of a domestically produced and an imported variety. At the foot of the tree the domestic variety is drawn from a single national market, a point taken up in the following subsection.

Production in industry and services is subject to an external economy of scale of the Marshallian type. The productivity term entering the value-added function is the product of a constant and the ratio of current to benchmark output in the same sector and the same state, raised to an elasticity that is set to zero in agriculture and to small positive values in industry and services. Because the term depends upon sectoral output in the state rather than upon the output of the individual firm, it is external to the firm and taken as given in the cost-minimisation problem. Competitive pricing and the zero pure profit condition are therefore preserved: the price of output is the value of value added plus the value of intermediate consumption per unit of output, with the value of value added exhausted by payments to labour and capital. The purchaser price of the output of a sector adds an indirect production tax margin to that unit cost, and the cost of labour to the employer is the wage received by the worker grossed up by a payroll levy, which is set to zero in the present calibration.

3.3 Output allocation, trade and the price system

Figure 2 sets out the treatment of trade, which is where the spatial structure of the model resides. The left-hand branch governs the disposal of output and the right-hand branch the sourcing of the composite commodity absorbed within a state.



Imports and exports each carry a state-specific ad valorem wedge on the international leg. The domestic leg carries no interstate wedge: sales into and purchases out of the national market take place at a single price common to every state.

Figure 2 Allocation of output and sourcing of the composite commodity

Producers allocate gross output between export to the rest of the world and sale on the domestic market through a constant-elasticity-of-transformation frontier, so that the export share of a state and sector rises with the price obtainable abroad relative to the price obtainable at home, at a rate governed by the elasticity of transformation. On the demand side, the composite commodity absorbed in a state is an Armington aggregate of the imported and the domestically supplied variety, so that the import share rises with the price of the domestic variety relative to the landed price of the import. The value of the composite equals the sum of the values of its two components, which defines the composite price faced by households and by producers using the commodity as an intermediate input.

Two features of the price system carry the spatial content of the model. The first is that both international legs are state specific. The price at which a state clears its exports is the world price converted at the nominal exchange rate and divided by the product of two wedges, one a state-specific ad valorem cost of reaching the world market and the other the export tax wedge that the simulation perturbs. The landed price of an import is correspondingly the world price converted at the exchange rate and marked up by the tariff and by the same state-specific wedge, which is applied symmetrically in the two directions. It is this wedge that represents the differential access of states to international markets, and it is the structural source of the coastal advantage discussed in Section 2.3.

The second feature is that the domestic leg carries no such wedge. The price at which a producer sells into the domestic market is a single national price for each commodity, identical across states, and the market for the domestic commodity clears at the national rather than the state level, the sum of domestic sales across states equalling the sum of domestic absorption across them. Purchaser prices for the domestic variety differ across states only through a commodity- and state-specific indirect tax margin. The consequence is that a change in the price of a domestically produced good reaches every state in full and without attenuation, which is the assumption examined at length in Section 2.5 and which the sensitivity analysis proposed there is designed to relax. World prices are held fixed, the country being treated as small in world markets, while the nominal exchange rate is endogenous.

3.4 Household demand and the housing market

Each state contains a single representative household whose preferences are represented by the nesting structure in Figure 3.

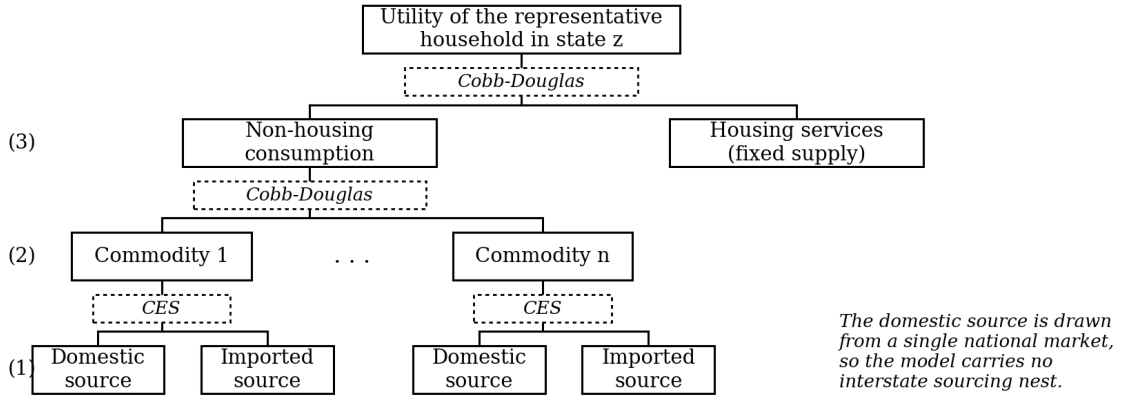


Figure 3 Nesting structure of state household demand

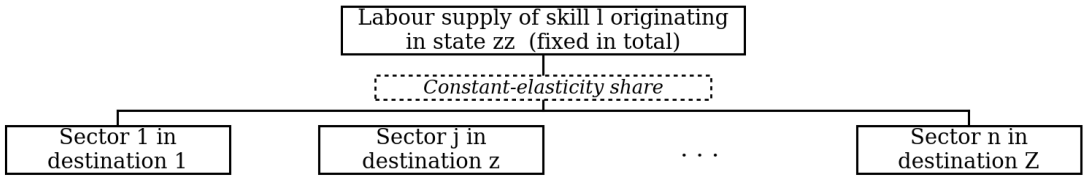
Household income is the sum of labour and capital earnings. Labour income is the wage bill accruing to workers employed within the state, summed over the origins from which those workers are drawn, so that a worker who moves is treated as a household that has relocated and whose income, expenditure and housing demand are recorded at the destination. Capital income is the return on the capital stock located in the state. Disposable income is income net of a direct tax that is fixed at its benchmark value, savings are a fixed proportion of disposable income, and the remainder is total consumption expenditure.

That expenditure is allocated at the top level between housing services and non-housing consumption in fixed value shares, the housing share being taken from benchmark expenditure data for each state. Non-housing expenditure is in turn allocated across the three composite commodities in fixed value shares, so that demand for each commodity is the product of the state consumption share, the non-housing share of expenditure and total consumption, divided by the composite price. Each commodity is sourced through the Armington nest already described. Because the housing component of the benchmark data is embedded in consumption of services, it is carved out of that item before the non-housing shares are calibrated, so that housing is not counted twice.

Housing is supplied inelastically. The stock available in each state is calibrated as benchmark housing expenditure valued at a unit benchmark rent, and rent adjusts so that its product with the fixed stock equals the housing share of total consumption expenditure in the state. A state that attracts workers therefore experiences a rise in rent that raises its cost of living and, through the migration equation described next, damps the inflow. Where the benchmark data record no housing expenditure the rent is held at its benchmark value and the market equation is suppressed. Government consumption, investment and changes in stocks are held at their benchmark levels throughout, so the model contains no investment behaviour and no government fiscal response.

3.5 Labour supply and interstate allocation

The treatment of labour is the principal departure from the standard interregional template and is set out in Figure 4. The unit of account is an arc, defined by a skill, a sector, a state of origin and a state of destination.



Every arc is valued at the destination wage, deflated by the state cost-of-living index and by housing rent, and net of a route-specific mobility wedge. Shares respond to that real return with a constant elasticity, so real returns are not equalised across arcs in equilibrium. Labour income accrues to the state of employment rather than to the state of origin.

Figure 4 Allocation of labour across sectors and states

The total supply of each skill originating in each state is fixed, so that the model contains no participation margin and no aggregate labour supply response; what is determined is the allocation of a given pool across sectors and destinations. Each arc is valued at a real return equal to the wage paid in the destination sector, deflated by the state cost-of-living index and by housing rent in the shares that govern the consumption nest, and divided by a route-specific mobility wedge that represents the cost of moving from the origin to the destination. The share of the origin pool allocated to a given arc is proportional to a route-specific constant multiplied by that real return raised to a common elasticity, normalised across all open arcs, which is the multinomial share form

discussed in Section 2.2. Employment of a skill in a sector and state is the sum of the arc flows terminating there, and this is what clears the labour market against the derived demand from the production nest.

The route-specific constants are calibrated to reproduce the observed base-year allocation at the observed base-year real returns. Two implications follow and should be stated plainly. Real returns are not equalised across arcs at the solution, since the elasticity is finite and the base-year differentials are absorbed into the constants; and the constants themselves do not respond to the shock, so that liberalisation moves the allocation along the calibrated share schedule rather than altering the schedule. The set of feasible arcs is restricted in the baseline to those on which a flow is observed in the base year. A variant that opens the full network is reported among the sensitivity results, and since the calibration assigns no weight to arcs with no observed flow, that variant requires a seed value to be imposed on the newly opened arcs. The seed has no counterpart in the data and is treated accordingly.

Capital is treated more simply. The stock in each state is fixed, it is mobile across sectors within the state, and a single rental rate clears the state capital market. There is thus no interstate reallocation of capital and no capital accumulation, which is the appropriate restriction for a comparative-static exercise and which places the burden of adjustment on labour reallocation and on relative prices.

3.6 Market clearing, closure and numeraire

Absorption of each composite commodity in each state is the sum of household consumption, intermediate demand from all sectors in the state, and the exogenous government, investment and stock components. The capital market clears within each state and the labour market clears arc by arc as described above. The market for each domestic commodity clears nationally. The price of the agricultural commodity in one state serves as numeraire, so that all results are to be read as relative prices and relative quantities. World prices are fixed and the nominal exchange rate is free to adjust, with foreign savings defined as the value of imports less the value of exports at world prices. The model has no dynamic attributes and no capital accumulation relations, so every result reported below answers a comparative-static question: given the structure of the economy in the benchmark year, what allocation would obtain under the counterfactual policy, compared with the allocation that would obtain in its absence.

3.7 Structural database

Calibration requires state-level data on gross output and value added by sector, on intermediate flows, on the split of value added between labour and capital and between the two skill categories, on household consumption by commodity, on housing expenditure shares, on exports, imports and domestic absorption by state and commodity, and on the matrix of worker flows by skill, sector, origin and destination. The benchmark is a fully specified state-level social accounting matrix for 2011, supplemented by consumption expenditure shares from the household expenditure survey, which are used to construct the alternative deflator reported alongside the principal welfare measure.

Two elements of the factor-payment split are assumptions rather than observations and are recorded here for transparency. The division of value added between labour and capital uses sectoral labour compensation shares of 0.154 in agriculture, 0.232 in industry and 0.418 in services. The division of the resulting labour payment between the two skill categories assumes that a skilled worker earns twice what an unskilled worker earns in the same sector and state, and the benchmark wage of each type follows as the implied payment per worker. Neither assumption is estimated within the model, and the sensitivity of the results to the second is examined below.

3.8 Behavioural parameters

Direct econometric estimates of the substitution elasticities required by a model of this dimension are not available for Indian states, so the values used are drawn from the ranges conventional in the applied literature and are set to reflect the relative flexibility of the three sectors. Substitution between labour and capital in value added is set to 0.30 in agriculture, 0.40 in industry and 0.50 in services, and substitution between skilled and unskilled labour to 0.30, 0.50 and 0.60 across the same three sectors, the ordering in each case reflecting the greater scope for factor substitution in activities less tied to land. On the trade side the Armington elasticity between imported and domestic varieties is set to 1.50 in agriculture, 2.00 in industry and 1.10 in services, and the elasticity of transformation between export and domestic sales to 1.20, 2.00 and 0.75, the low value for services reflecting the limited tradability of much of that sector. The elasticity governing the allocation of workers across arcs is set to 4 for both skill categories. The scale elasticity is set to zero in agriculture, 0.02 in industry and 0.01 in services,

magnitudes small enough to satisfy the requirement, discussed in Section 2.4, that agglomeration forces be bounded relative to the dispersion forces if the equilibrium is to be well behaved.

3.9 The simulated shock and the welfare decomposition

The counterfactual is a 20 per cent reduction in the export-tax wedge on every commodity, applied uniformly across states and sectors, with all other policy instruments held at their benchmark values. Because the wedge enters the export price of every state divided by the same state-specific access cost, a uniform national liberalisation delivers a differentiated effective shock across states, and it is that differentiation, interacting with differences in sectoral composition, in factor endowments and in consumption baskets, that generates the regional pattern reported in the following section.

The welfare measure is real household consumption in each state, defined as total nominal consumption expenditure deflated by a price index that combines the state consumption price index over non-housing commodities with housing rent, weighted by the shares governing the consumption nest. Written in logarithms, the change in real consumption is the change in nominal consumption less the change in that price index, and the two terms are the output-led and the price-led components of the decomposition on which the analysis rests. The first records what happens to a state as a producer, through the value of what its factors earn; the second records what happens to it as a consumer, through the cost of the basket its households buy. The decomposition is exact by construction rather than approximate, and it requires no assumption beyond those already embedded in the consumption nest. A second version of the measure, in which the deflator is constructed from household expenditure survey shares rather than from the social accounting matrix, is reported alongside the first as a check that the pattern is not an artifact of the benchmark weights.

4. Results

4.1 The aggregate effect and the three-way regional pattern

The counterfactual reduces the export-tax wedge on every commodity by 20 per cent, uniformly across states and sectors. The direct effect is a rise in the price received by exporters, and the immediate response is an

expansion of export volumes in every zone, by 1.22 per cent on the coast, 0.92 per cent in the hill states and 0.68 per cent in the interior. Imports rise by a comparable amount everywhere, between 1.08 and 1.30 per cent, as the exchange rate adjusts and the additional foreign exchange is spent. Real household consumption rises in all three zones. The interest of the exercise lies not in that aggregate result but in its distribution, and in the fact that the ordering of the zones by welfare gain cannot be recovered from the ordering of their production responses.

Table 1 reports the central result. All figures are percentage changes from the benchmark, aggregated to the zone level using state gross output as weights.

Table 1 Welfare change and its decomposition, by zone (percentage change)

Zone	Output-led	Price-led	Total real	Price-led share
Coastal	0.049	0.193	0.243	80%
Hilly	0.004	0.208	0.212	98%
Interior	-0.057	0.162	0.105	154%

The three zones fall into a clear order, the coastal states gaining most at 0.243 per cent, the hill states close behind at 0.212 per cent and the interior states gaining least at 0.105 per cent, less than half the coastal figure. That ordering is the expected one and would be predicted by any account resting on market access. What is not predicted is the composition of the three gains, which differs so sharply across the zones that the same total can be reached by different routes and a common route can produce different totals.

4.2 The decomposition of welfare change

The welfare change in each state has been separated into an output-led component, the change in nominal consumption financed by what the state earns as a producer, and a price-led component, the fall in the cost of the basket its households buy. Figure 5 sets the two against the total.

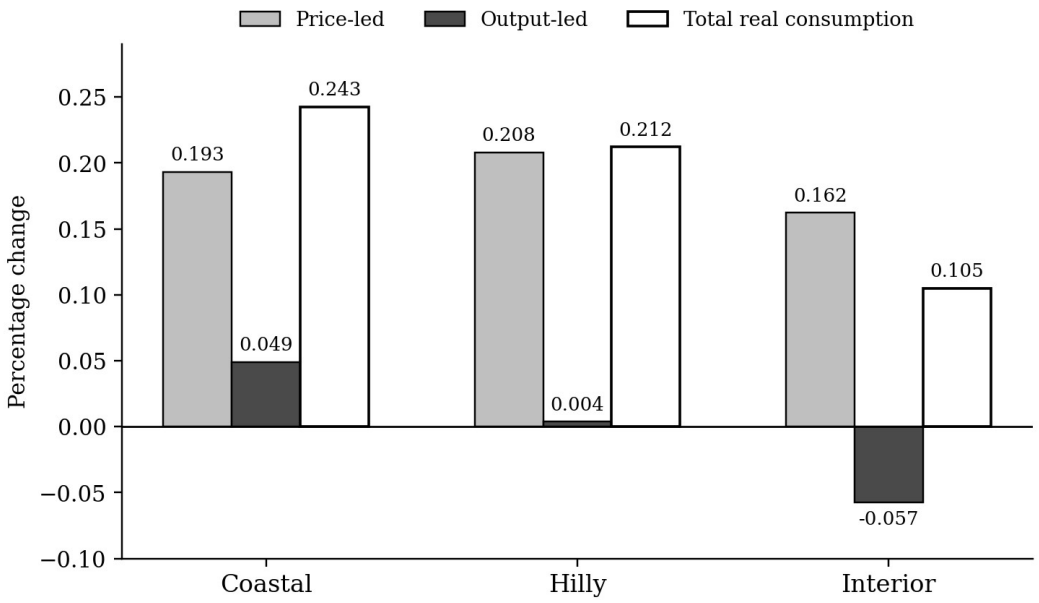


Figure 5 Output-led and price-led components of welfare change, by zone

Two features stand out. The first is that the price-led component is large and positive in every zone, and that it is the larger of the two components everywhere. Liberalisation lowers the price of traded goods, and because those goods enter every state consumption basket, every state gains as a consumer irrespective of what happens to its production. Between 80 and 154 per cent of each zone gain is accounted for by this channel. The second feature is that the output-led component, though smaller in magnitude, is the only one that changes sign. It is positive on the coast at 0.049 per cent, almost exactly zero in the hill states at 0.004 per cent, and negative in the interior at minus 0.057 per cent. The consumption channel therefore establishes the level of the gains, which is broadly common, while the production channel establishes the differences between them.

This is the sense in which the sign of a state outcome cannot be read from its production structure alone. The interior states are worse off as producers: their nominal consumption falls, their output contracts by 0.110 per cent and the earnings of their factors decline. Considered on the production side alone, they are losers from the reform. They nevertheless record a welfare gain, because the fall in the price of what they buy is more than sufficient to offset the fall in what they earn, the price-led component standing at 154 per cent of the net outcome. An analysis conducted in terms of output, factor earnings or sectoral reallocation, which is how the regional

incidence of trade reform is usually measured, would record the interior as suffering a loss and would miss the offset entirely.

The hill states present the mirror image and are the more striking case. Their output-led component is 0.004 per cent, which is to say that the production side of their economies contributes nothing at all to their welfare outcome. Their gain of 0.212 per cent, second only to the coast, is 98 per cent price-led. It arises because their consumption baskets are oriented towards industrial goods, the category whose price falls furthest under liberalisation, so that the same national price change delivers a larger reduction in their cost of living than in that of any other zone: their consumption price index falls by 0.210 per cent against 0.195 on the coast and 0.163 in the interior. A zone remote from ports, whose output contracts, and whose factor earnings barely move, nevertheless finishes close to the top of the welfare ranking on the strength of what its households consume.

4.3 The production response and the coastal advantage

Table 2 sets out the production side, which is where the zones diverge.

Table 2 Production, trade and factor market outcomes, by zone (percentage change)

Zone	Output	Exports	Imports	Labour demand	Labour income
Coastal	0.132	1.220	1.297	0.000	-0.065
Hilly	-0.073	0.917	1.106	0.070	0.000
Interior	-0.110	0.680	1.080	-0.005	-0.086

Note: The labour-demand column is calculated from summed benchmark and counterfactual employment levels. It is not an average of state percentage changes.

The coastal states are the only zone in which output expands, by 0.132 per cent, and they record the largest export response at 1.220 per cent. The mechanism is the one set out in Section 3.3. Because the export price received by a state is the world price divided by a state-specific access wedge, a uniform national reduction in the export tax delivers a larger effective price improvement to states that are cheaply connected to world markets, and the resources of those states are drawn into traded production. The mild increasing returns in industry reinforce the movement, since a state whose industrial output expands becomes more productive and expands further. Both mechanisms operate in the same direction and both favour the coast.

The interior states show the opposite configuration. Their export response is the weakest of the three at 0.680 per cent, their output contracts by 0.110 per cent, and their labour income falls by 0.086 per cent. Their production is concentrated in agriculture and in non-traded activity, sectors which face the lowest transformation and substitution elasticities in the model and which therefore respond least to a change in export incentives, while their remoteness from ports means that the price improvement reaching them is in any case smaller. Resources cannot leave the sectors in which they are trapped quickly enough, nor the states in which they are located cheaply enough, to take advantage of the expansion occurring elsewhere. The incidence of a common national shock is thus converted into a durable regional imprint, which is the mechanism that the literature reviewed in Section 2.2 leads one to expect.

4.4 Factor market incidence and the skill dimension

Table 3 disaggregates the wage response by skill, and it is here that the incidence of the reform is sharpest.

Table 3 Real wage change by skill and zone (percentage change)

Zone	Real wage, skilled	Real wage, unskilled	Nominal wage, skilled	Nominal wage, unskilled
Coastal	0.142	0.132	-0.052	-0.062
Hilly	0.031	0.154	-0.177	-0.055
Interior	-0.174	0.111	-0.336	-0.051

Unskilled real wages rise in every zone and by similar amounts, between 0.111 and 0.154 per cent, so that the geography of the reform is barely visible in the return to unskilled labour. Skilled wages behave quite differently. Skilled real wages rise by 0.142 per cent on the coast, rise slightly in the hill states, and fall by 0.174 per cent in the interior, the only real wage decline recorded anywhere in the exercise. The nominal figures show the same pattern more starkly: skilled nominal wages fall by 0.336 per cent in the interior against 0.052 per cent on the coast. The regional incidence of the reform is therefore concentrated almost entirely in the market for skilled labour.

The employment decomposition is consistent with weaker demand for skilled labour in the interior, but it does not by itself identify migration routes. Skilled destination employment rises by 0.016 per cent in the coastal

states and by 0.084 per cent in the hilly states, while declining by 0.030 per cent in the interior. Unskilled employment is effectively unchanged on the coast and in the interior and rises modestly in the hills. Because destination employment sums workers from all states of origin, these figures identify where employment is located after the shock rather than whether workers have crossed state borders. The wage pattern should therefore be read as evidence of differentiated regional demand for skill, not as a direct estimate of interstate migration.

4.5 Labour reallocation

The labour-market closure fixes the total supply of each skill category by state of origin. The counterfactual therefore does not create or destroy employment nationally; it reallocates a fixed labour pool across sectors and destination states. Reconstructing counterfactual employment from the benchmark labour levels and the reported percentage changes confirms this restriction: total skilled and unskilled employment are unchanged up to rounding.

Regional employment changes are calculated by summing benchmark and counterfactual employment levels before taking percentage changes. On this basis, total employment is essentially unchanged in the coastal states, falling by only 0.0004 per cent; it rises by 0.0698 per cent in the hilly states and falls by 0.0051 per cent in the interior. These level-consistent estimates are smaller than averages of state growth rates and show that the aggregate labour response is modest.

The sectoral decomposition is more informative than the regional total. In the coastal states, agricultural employment falls by 0.118 per cent, while manufacturing and service employment rise by 0.346 and 0.029 per cent respectively. Coastal output expansion is therefore accompanied by a shift of employment out of agriculture and into manufacturing and services, rather than by a contraction of manufacturing employment. In the hilly states, the 0.070 per cent increase in total employment is concentrated in agriculture, where employment rises by 0.196 per cent, while manufacturing and services decline by 0.279 and 0.043 per cent. In the interior, agricultural employment rises by 0.159 per cent but is offset by contractions of 0.355 per cent in manufacturing and 0.252 per cent in services.

These results establish sectoral and spatial reallocation but not interstate migration. Destination employment includes workers whose state of origin and destination are the same, so a direct migration interpretation would

require the origin-destination labour-flow matrix. The defensible conclusion is narrower: the shock produces little change in total regional employment but a marked change in its sectoral composition, with manufacturing employment expanding on the coast and contracting in both non-coastal groups.

The magnitudes throughout are small. A 20 per cent reduction in the export-tax wedge moves real consumption by roughly a quarter of one per cent in the most favoured zone, while the resulting changes in total regional employment remain below one tenth of one per cent. The exercise is therefore most informative about the direction, composition and relative regional incidence of adjustment rather than its absolute scale.

4.6 Cross-state manufacturing employment response

The sectoral decomposition suggests that manufacturing employment is the margin on which the regional divide is sharpest. To examine this pattern across the twenty-four states, the percentage change in state manufacturing employment is regressed on the benchmark manufacturing-employment share, the benchmark skilled-employment share and regional indicators for hilly and interior states, with coastal states as the omitted category. The model is weighted by benchmark manufacturing employment so that states with larger manufacturing labour forces receive greater economic weight, and heteroskedasticity-robust standard errors are reported. The estimates are descriptive associations among simulated outcomes and benchmark characteristics; they are not causal estimates.

The weighted regression is jointly significant, with $F(4,19) = 17.38$ and an R-squared of 0.651. Neither the benchmark manufacturing-employment share nor the skilled-employment share is statistically significant. Regional location, however, is strongly associated with the response. The regional indicators are jointly significant, $F(2,19) = 15.34$, $p = 0.0001$. Relative to coastal states, the manufacturing-employment response is 0.661 percentage points lower in hilly states and 0.716 percentage points lower in interior states.

Table 4 Cross-state correlates of manufacturing employment change

Variable	Estimate	Robust SE	p-value
Panel A. Weighted regression			
Manufacturing employment share	-0.0047	0.0249	0.851
Skilled employment share	0.0041	0.0162	0.801

Hilly state indicator	-0.6606	0.2351	0.011
Interior state indicator	-0.7158	0.1382	<0.001
Constant	0.3858	0.2876	0.196
Panel B. Adjusted regional margins			
Coastal	0.3521	0.0541	<0.001
Hilly	-0.3085	0.2067	0.152
Interior	-0.3638	0.1309	0.012

Notes: Dependent variable is the percentage change in manufacturing employment. Analytic weights are benchmark manufacturing employment. Coastal states are the reference group. N = 24; R-squared = 0.6509. Robust standard errors are reported.

The adjusted regional margins sharpen the result. Manufacturing employment rises by 0.352 per cent in coastal states but falls by 0.308 per cent in hilly states and 0.364 per cent in interior states. After a Bonferroni correction for the three pairwise comparisons, the hilly-coastal difference remains significant at $p = 0.034$ and the interior-coastal difference at $p < 0.001$. The hilly and interior responses are statistically indistinguishable, with a difference of only 0.055 percentage points and an adjusted p-value of 1.000. The relevant empirical divide is therefore coastal versus non-coastal, rather than hilly versus interior.

The absence of significant coefficients on the initial manufacturing and skill shares implies that these two benchmark characteristics do not explain the regional gap. The regional indicators capture a broader bundle of location-specific conditions, potentially including access to international markets, infrastructure, industrial linkages and agglomeration forces. Because those mechanisms are not separately identified in this regression, the result should be stated as evidence of a robust coastal association rather than as a causal estimate of port access. An unweighted specification produces the same ordering, although the hilly-coastal contrast is less precisely estimated, which indicates that the weighted result is not driven solely by small coastal states.

4.7 Robustness and limitations

Three checks bear on the reliability of the pattern reported above. The first concerns the deflator. The welfare measure reported here uses a price index constructed from the benchmark social accounting matrix; recomputing it from household expenditure survey shares changes the level of the estimated gains slightly but

leaves the ordering of the three zones and the composition of each gain unaltered. The second concerns housing. Because housing is in fixed supply and rent is determined by the housing share of consumption expenditure, rent moves in proportion to nominal consumption in every state, and the difference between a deflator that includes rent and one that excludes it is of the order of two thousandths of a percentage point. Housing therefore operates in this model as a component of the cost of living rather than as an independent congestion mechanism damping migration, and no part of the regional pattern reported above depends upon it.

The third and most substantial qualification concerns the composition of the hilly zone. Two states, Uttaranchal and Himachal Pradesh, account for approximately two thirds of the gross output of that zone and therefore dominate its weighted averages. Both have industrial shares far above what their geography would suggest, a consequence of the central excise exemptions extended to those states in 2003, which drew pharmaceutical, consumer goods and automotive component manufacturing into locations that would otherwise have remained agricultural. The industrial orientation of the hilly consumption basket, which is what drives the price-led result, is thus in part a fiscal artifact rather than a structural feature of mountainous economies. This does not weaken the analytical point, since the mechanism operates through the composition of the basket whatever its historical origin, but it does limit the generality of the hilly finding, which should be read as a statement about these particular states rather than about hill economies at large.

5. Conclusion

This paper has examined the regional incidence of a uniform 20 per cent reduction in India's export-tax wedge using a multi-region computable general equilibrium model calibrated to twenty-four states. The main result is that a national trade reform does not produce a spatially uniform outcome. Coastal states obtain the largest welfare and production gains, the interior experiences a contraction in output and factor earnings, and the hilly states remain comparatively resilient despite a weak production response. The common policy shock is filtered through state differences in international market access, production structure, factor adjustment and household expenditure patterns.

The welfare decomposition is central to this conclusion. Real consumption rises in every region, but the dominant contribution comes from lower consumer prices rather than higher nominal income. The interior states lose on the production side yet retain a positive welfare gain because the fall in their cost of living more than offsets the decline in what they earn. The hilly states provide an even clearer illustration: their production contribution is close to zero, while nearly the entire welfare gain is generated by the price channel. Regional welfare therefore cannot be inferred from output alone. What states consume is as important as what they produce.

The labour results refine the production-side story. Since labour supply is fixed by skill and state of origin, the model generates reallocation rather than national employment creation. Total regional employment changes very little, but its sectoral composition changes materially. Coastal agricultural employment contracts while manufacturing and service employment expand. The hilly increase in employment is concentrated in agriculture, and the interior combines agricultural expansion with contractions in manufacturing and services. The cross-state regression confirms that manufacturing employment is the margin on which the coastal advantage is most pronounced: adjusted manufacturing employment rises on the coast and falls in both non-coastal groups, with the coastal contrasts remaining significant after correction for multiple comparisons.

At the same time, the regression does not show that the coastal advantage is explained by a larger initial manufacturing base or a higher skilled-employment share. Those variables are statistically insignificant once regional location is included. The regional indicators therefore represent a wider set of structural conditions that the present cross-sectional exercise cannot separate. The finding is consistent with a role for market access, infrastructure, industrial linkages and agglomeration economies, but it should not be interpreted as a causal test of any one mechanism.

The policy implication is that place-neutral liberalisation can generate place-specific adjustment. Trade reform may improve aggregate and household welfare while simultaneously weakening production and manufacturing employment in regions that are poorly connected to external markets. Complementary policies should therefore focus on reducing domestic connectivity costs, strengthening non-coastal industrial linkages, supporting sectoral adjustment and enabling workers to move without losing access to housing and public services. Such measures would not replace trade reform; they would determine how widely its production-side gains are shared.

The analysis remains subject to important limitations. It is comparative static, uses a 2011 benchmark, imposes common behavioural elasticities, treats the domestic commodity market as nationally integrated without explicit internal transport costs, and represents each state through a single household. The hilly result is also influenced by the unusual industrial structure of Himachal Pradesh and Uttaranchal. Finally, the cross-state regressions use simulated outcomes and only twenty-four observations. The estimates are best viewed as structured descriptive evidence that complements the model rather than as independent econometric identification. Future work could introduce bilateral domestic trade costs, distinguish household groups within states, and use the full origin-destination labour matrix to separate sectoral reallocation from interstate migration.

Acknowledgements

The author received no financial support for the research, authorship, or publication of this paper and declares no financial relationships or other potential conflicts of interest.

References

Allen, T. and Arkolakis, C. (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics*, 129(3), 1085-1140.

Artuc, E., Chaudhuri, S. and McLaren, J. (2010). Trade shocks and labor adjustment: A structural empirical approach. *American Economic Review*, 100(3), 1008-1045.

Atkin, D. and Donaldson, D. (2015). Who's getting globalized? The size and implications of intra-national trade costs. NBER Working Paper No. 21439. Cambridge, MA: National Bureau of Economic Research.

Attanasio, O., Goldberg, P. K. and Pavcnik, N. (2004). Trade reforms and wage inequality in Colombia. *Journal of Development Economics*, 74(2), 331-366.

Autor, D. H., Dorn, D. and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6), 2121-2168.

Barua, A. and Chakraborty, P. (2010). Does openness affect regional inequality? A case study for India. *Review of Development Economics*, 14(3), 447-465.

Boeters, S. and Savard, L. (2013). The labour market in CGE models. In P. B. Dixon and D. W. Jorgenson (eds.), *Handbook of Computable General Equilibrium Modeling*, Vol. 1B (pp. 1645-1718). Amsterdam: Elsevier.

Bussolo, M. and Round, J. I. (eds.) (2006). *Globalisation and Poverty: Channels and Policy Responses*. London: Routledge.

Caliendo, L., Dvorkin, M. and Parro, F. (2019). Trade and labor market dynamics: General equilibrium analysis of the China trade shock. *Econometrica*, 87(3), 741-835.

Caliendo, L., Parro, F., Rossi-Hansberg, E. and Sarte, P.-D. (2018). The impact of regional and sectoral productivity changes on the U.S. economy. *The Review of Economic Studies*, 85(4), 2042-2096.

Chen, S. and Ravallion, M. (2004). Welfare impacts of China's accession to the World Trade Organization. *The World Bank Economic Review*, 18(1), 29-57.

Cockburn, J., Decaluwe, B. and Robichaud, V. (eds.) (2008). *Trade Liberalization and Poverty: A CGE Analysis of the 1990s Experience in Africa and Asia*. Poverty and Economic Policy (PEP) Research Network.

Cosar, A. K. and Fajgelbaum, P. D. (2016). Internal geography, international trade, and regional specialization. *American Economic Journal: Microeconomics*, 8(1), 24-56.

Daumal, M. (2013). The impact of trade openness on regional inequality: The cases of India and Brazil. *The International Trade Journal*, 27(3), 243-280.

Dix-Carneiro, R. and Kovak, B. K. (2017). Trade liberalization and regional dynamics. *American Economic Review*, 107(10), 2908-2946.

Dixon, P. B. and Jorgenson, D. W. (eds.) (2013). *Handbook of Computable General Equilibrium Modeling*, Vols. 1A and 1B. Amsterdam: Elsevier.

Donaldson, D. (2018). Railroads of the Raj: Estimating the impact of transportation infrastructure. *American Economic Review*, 108(4-5), 899-934.

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5), 1741-1779.

Fajgelbaum, P. D. and Khandelwal, A. K. (2016). Measuring the unequal gains from trade. *The Quarterly Journal of Economics*, 131(3), 1113-1180.

Fujita, M., Krugman, P. and Venables, A. J. (1999). *The Spatial Economy: Cities, Regions, and International Trade*. Cambridge, MA: MIT Press.

Galrao Carneiro, F. and Arbache, J. S. (2003). The impacts of trade on the Brazilian labor market: A CGE model approach. *World Development*, 31(9), 1581-1595.

Haddad, E. A. and Hewings, G. J. D. (1998). *Transportation costs and regional development: An interregional CGE analysis*. Paper presented at the 38th European Congress of the Regional Science Association, Vienna.

Haddad, E. A. and Hewings, G. J. D. (2004). *Transportation costs, increasing returns and regional growth: An interregional CGE analysis*. Working Paper 315, Department of Economics, University of Brasilia.

Haddad, E. A. and Hewings, G. J. D. (2005). Market imperfections in a spatial economy: Some experimental results. *The Quarterly Review of Economics and Finance*, 45(2-3), 476-496.

Haddad, E. A., Domingues, E. P. and Perobelli, F. S. (2002). Regional effects of economic integration: The case of Brazil. *Journal of Policy Modeling*, 24(5), 453-482.

Haddad, E. A., Hewings, G. J. D., Perobelli, F. S. and dos Santos, R. A. C. (2010). Regional effects of port infrastructure: A spatial CGE application to Brazil. *International Regional Science Review*, 33(3), 239-263.

Haddad, E. A., Lahr, M. L., Elshahawany, D. N. and Vassallo, M. (2016). Regional analysis of domestic integration in Egypt: An interregional CGE approach. *Journal of Economic Structures*, 5(25), 1-24.

Hanson, G. H. (2005). Market potential, increasing returns and geographic concentration. *Journal of International Economics*, 67(1), 1-24.

Harris, C. D. (1954). The market as a factor in the localization of industry in the United States. *Annals of the Association of American Geographers*, 44(4), 315-348.

Hasan, R., Mitra, D., Ranjan, P. and Ahsan, R. N. (2012). Trade liberalization and unemployment: Theory and evidence from India. *Journal of Development Economics*, 97(2), 269-280.

Hinojosa-Ojeda, R. A., Robinson, S. and De Paolis, F. (1999). Regional integration among the unequal: A CGE model of NAFTA and the Central American republics. *The North American Journal of Economics and Finance*, 10(1), 235-292.

Horridge, M. (2011). *The TERM model and its database.* General Paper G-219, Centre of Policy Studies, Monash University.

Horridge, M., Yusuf, A. A., Ginting, E. and Aji, P. (2016). *Improving Indonesia's domestic connectivity: An inter-regional CGE analysis*. ADB Papers on Indonesia No. 17. Manila: Asian Development Bank.

Jones, R. and Whalley, J. (1989). A Canadian regional general equilibrium model and some applications. *Journal of Urban Economics*, 25(3), 368-404.

Kambourov, G. (2009). Labour market regulations and the sectoral reallocation of workers: The case of trade reforms. *Review of Economic Studies*, 76(4), 1321-1358.

Kennan, J. and Walker, J. R. (2011). The effect of expected income on individual migration decisions. *Econometrica*, 79(1), 211-251.

Kovak, B. K. (2013). Regional effects of trade reform: What is the correct measure of liberalization? *American Economic Review*, 103(5), 1960-1976.

Krugman, P. (1991). Increasing returns and economic geography. *Journal of Political Economy*, 99(3), 483-499.

Mishra, P. and Kumar, U. (2005). *Trade liberalization and wage inequality: Evidence from India*. IMF Working Paper WP/05/20. Washington, DC: International Monetary Fund.

Naranpanawa, A. and Arora, R. (2014). Does trade liberalization promote regional disparities? Evidence from a multiregional CGE model of India. *World Development*, 64, 339-349.

Paluzie, E., Pons, J. and Tirado, D. A. (2004). The geographical concentration of industry across Spanish regions, 1856-1995. *Jahrbuch für Regionalwissenschaft (Review of Regional Research)*, 24(2), 143-160.

Porto, G. G. (2006). Using survey data to assess the distributional effects of trade policy. *Journal of International Economics*, 70(1), 140-160.

Pradhan, B. K. and Sahoo, A. (2008). The impact of trade liberalization on household welfare and poverty in India. In J. Cockburn, B. Decaluwe and V. Robichaud (eds.), *Trade Liberalization and Poverty: A CGE Analysis of the 1990s Experience in Africa and Asia* (pp. 217-262). Poverty and Economic Policy (PEP) Research Network.

Redding, S. J. (2016). Goods trade, factor mobility and welfare. *Journal of International Economics*, 101, 148-167.

Redding, S. J. (2020). *Trade and geography*. CEP Discussion Paper No. 1718. London: Centre for Economic Performance, LSE.

Redding, S. J. and Rossi-Hansberg, E. (2017). Quantitative spatial economics. *Annual Review of Economics*, 9, 21-58.

Redding, S. and Venables, A. J. (2004). Economic geography and international inequality. *Journal of International Economics*, 62(1), 53-82.

Rodriguez-Pose, A. (2010). *Trade and regional inequality*. Policy Research Working Paper 5347. Washington, DC: World Bank.

Rodriguez-Pose, A. and Gill, N. (2004). Is there a global link between regional disparities and devolution? *Environment and Planning A*, 36(12), 2097-2117.

Rodriguez-Pose, A. and Gill, N. (2006). How does trade affect regional disparities? *World Development*, 34(7), 1201-1222.

Sanyal, S. and Arora, A. (2024). *Relative economic performance of Indian states: 1960-61 to 2023-24*. EAC-PM Working Paper Series EAC-PM/WP/31/2024. New Delhi: Economic Advisory Council to the Prime Minister.

Sinha, A. and Adam, C. (2004). *Trade reforms and informal economy in India: A CGE analysis*. National Council of Applied Economic Research, New Delhi, and Queen Elizabeth House, University of Oxford.

Timilsina, G. R., Sahoo, P. and Dash, R. K. (2022). *Why do Indian states differ in their infrastructure development?* Policy Research Working Paper 10086. Washington, DC: World Bank.

Topalova, P. (2007). Trade liberalization, poverty and inequality: Evidence from Indian districts. In A. Harrison (ed.), *Globalization and Poverty* (pp. 291-336). Chicago: University of Chicago Press / NBER.

Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India. *American Economic Journal: Applied Economics*, 2(4), 1-41.

Uneven benefits from trade liberalization: Tariff barriers and labor migration in China (2026). *Economic Modelling*. Advance online publication. [author names to be confirmed]

Ural Marchand, B. (2012). Tariff pass-through and the distributional effects of trade liberalization. *Journal of Development Economics*, 99(2), 265-281.

Van Leemput, E. (2021). A passage to India: Quantifying internal and external barriers to trade. *Journal of International Economics*, 131, 103473.

Wan, G., Lu, M. and Chen, Z. (2007). Globalization and regional income inequality: Empirical evidence from within China. *Review of Income and Wealth*, 53(1), 35-59.

Winters, L. A., McCulloch, N. and McKay, A. (2004). Trade liberalization and poverty: The evidence so far. *Journal of Economic Literature*, 42(1), 72-115.