MIT 14.76/760: Firms, Markets, Trade and Growth
Sp 2026, Lecture 18: Industrial and Place-Based Policy

Dave Donaldson

Industrial Policy and Place-Based Policy

- Basic idea: a deliberate attempt to use policy to promote the success (or prevent demise) of a particular, targeted industry or location

- Why might it make sense do this?

Industrial Policy and Place-Based Policy

- Basic idea: a deliberate attempt to use policy to promote the success (or prevent demise) of a particular, targeted industry or location

- Why might it make sense do this?

- As usual, two broad possible policy rationales:
  1. Pareto Efficiency:
     - For whatever reason, the targeted industry/location has higher VMPK and/or VMPL than other industries/locations
     - So aggregate productivity would rise if mobile inputs (K, L) were to move to the targeted industry/location
  2. Distributional concerns:
     - For whatever reason, policymaker prefers (people whose earnings are relatively tied to) a given industry/location
     - E.g. could derive from political realities (industries lobby; locations drive elections) or income-based redistribution motive

- Main focus today will be #1 as it gels with the themes of this class (#2 may be just as important, but also likely to be more context-specific)

Possible reasons for VMPL (or VMPK) differences

- Why might it be the case that the targeted industry/location has higher VMPL than other industries/locations?

Possible reasons for VMPL (or VMPK) differences

- Why might it be the case that the targeted industry/location has higher VMPL than other industries/locations?
- Recall, generically the answer would have to lie in *market failures*
- Without such market failures, firms' own private incentives would act to equalize VMPL across firms (and hence also industries/locations).
    - ⇒ In such a world, there is no such thing as an industry/location that we should deem as any "better" than some other...all industries/locations producing the same amount of social value of output (on the margin).
    - ⇒ "It doesn't make any difference whether a country makes computer chips or potato chips [on the margin]. $100 or $100 of the other, it's still $100." (Michael Boskin, 1992, as member of Pres. Bush's economic council)
    - ⇒ In such a world the best industrial policy is none at all
- Obviously the real world has market failures. The challenge in practice, however, is knowing where they are, and where they are relatively worse.

Possible reasons for VMPL (or VMPK) differences

- But what are some specific examples of market failures that we might worry about here (in the sense that they might plausibly differ across industries/locations)?

Possible reasons for VMPL (or VMPK) differences

- But what are some specific examples of market failures that we might worry about here (in the sense that they might plausibly differ across industries/locations)?

- Some *static* candidates:
  - Productivity spillovers across firms within industry/location; e.g.:
    - Knowledge embodied in workers who move across firms
    - Many things can't be patented
    - Evan patents disclose things useful for making next product
  - "Thick market effects" (lower search congestion externalities) for inputs (e.g. programmers or VC in Silicon Valley)
  - Imperfect competition—usually means firm doesn't have right social incentive (could be too high, or too low) to exploit internal RTS

- Some *dynamic* candidates:
  - Firms have (internal) increasing RTS but face borrowing constraints
  - Firms have (internal) learning-by-doing but face borrowing constraints
  - External learning-by-doing (or knowledge spillovers) at industry/location level

A Simple Model of Industrial/Place-Based Policy

- We will walk through an example of 2 "industries" (k), but could just as easily be "locations" or even "industry-locations"

- Suppose simple (and static) setting with:
  - Production by representative firm z in each industry is $y_k(z) = A_k l_k(z)$, where $l_k(z)$ is "labor" input by firm z
  - But—and this is the key—there is a potential *externality in production*:
    - $A_k = a_k L_k^{\gamma_k}$, where a_k is exogenous productivity and $L_k \equiv \sum_z l_k(z)$ is total labor used in k.
    - Each firm z takes L_k as given (firms correctly recognize that they are each too small for own actions to affect L_k much).
  - Total amount of labor available is $\bar{L}$. So $\bar{L} = L_1 + L_2$ always.
  - Labor earns wage w, competitive labor market
  - Consumers (= workers) have Cobb-Douglas preferences with weights α_k
  - Perfect competition in production ($p_k(z) = MC_k(z)$ will hold)

- So γ_k governs the strength of the production externality.

Laissez-Faire Market Allocation

- What will happen here?
- Perfect competition implies $p_k = w/A_k$.
  - Note how the definition of MC is $w \times [\frac{\partial y_k(z)}{\partial l_k(z)}]^{-1} = w/A_k$, i.e the derivative is from the firm's perspective, which holds L_k and hence A_k constant.
- Hence supply-side implies: $p_k y_k = wL_k$
- Cobb-Douglas demand implies $p_k y_k = \frac{\alpha_k}{\alpha_1+\alpha_2} Y = \frac{\alpha_k}{\alpha_1+\alpha_2}(w\bar{L})$
- So market equilibrium implies: $\frac{\alpha_k}{\alpha_1+\alpha_2} = \frac{L_k}{\bar{L}} \equiv \lambda_k$

First-Best Allocation

- As always, to solve for the social optimum, ask what a hypothetical "social planner" would do.

- Planner would choose (L_1, L_2) so as to maximize utility, subject to same technology ($y_k = A_k L_k$, with $A_k = L_k^{\gamma_k}$) and resource constraint ($\bar{L} = L_1 + L_2$) as market economy. That is, planner would solve:

$$\max_{\lambda_1} \ln U = \kappa + \alpha_1(1+\gamma_1)\ln(\lambda_1) + \alpha_2(1+\gamma_2)\ln(1-\lambda_1)$$

  where κ is some constant (involving a_1, a_2, and $\bar{L}$) that doesn't matter

- Solution to this problem (try it!) is:

$$\lambda_1^* = \frac{\alpha_1}{\alpha_2 \frac{1+\gamma_2}{1+\gamma_1} + \alpha_1}$$

First-Best Allocation—Interpretation

- What does this tell us?

- Compare the optimal allocation $\lambda_1^* = \frac{\alpha_1}{\alpha_2 \frac{1+\gamma_2}{1+\gamma_1} + \alpha_1}$ to the market allocation $\lambda_1 = \frac{\alpha_1}{\alpha_1 + \alpha_2}$

- Optimal allocation is same as market allocation (i.e. $\lambda_1 = \lambda_1^*$) iff $\gamma_1 = \gamma_2$.
  - That is, market allocation is efficient iff $\gamma_1 = \gamma_2$ (i.e. when externalities have equal elasticities in two sectors).

- Otherwise, whenever $\gamma_1 > \gamma_2$, optimal allocation would put more labor into sector 1 than the market does (ie $\lambda_1^* > \lambda_1$).
  - That is, the sector with the greater externality (elasticity) is too small in the market allocation
  - In the market allocation firms do not internalize the beneficial effects of their actions onto others in the same industry; the industry that does this least well is the one that is too small.

Optimal Industrial Policy

- Recall, competitive market allocation is inefficient when $\gamma_1 \neq \gamma_2$.
- Could "industrial policy" fix this problem?
- Trivial solution: just tell firms to produce such that $L_1 = \lambda_1^* \bar{L}$.
- More pragmatic solution: impose *production subsidies* s_k on revenue.
  - Firm's profits become: $\pi_k(z) = (1 + s_k) p_k y_k(z) - w l_k(z)$
  - WLOG, can set $s_2 = 0$ and solve for $s_1 > -1$ only.
  - NB: tax to pay the subsidy must be raised (lump-sum) from consumers, so demand will now satisfy $p_k y_k = \frac{\alpha_k}{\alpha_1 + \alpha_2}(w\bar{L} - s_1 p_1 y_1)$
- Market allocation (in presence of subsidy s_1) will now satisfy:

$$\lambda_1^s = \frac{\alpha_1}{\alpha_2 \frac{1}{(1+s_1)} + \alpha_1}$$

- So, comparing this to λ_1^*, we see that welfare-optimizing industrial policy (which we could call s_1^*) would be: $s_1^* = \frac{1+\gamma_1}{1+\gamma_2} - 1$

Practical matters

- What would this optimal subsidy look like in practice?

- Practical limitations:
  - We'd have to know how big γ_1/γ_2 is. Who would know? Would the firms? Even if they knew, would they tell the industrial policy-designer the truth?
  - How to verify revenues $(p_1 y_1(z))$ in order to pay out right subsidy $(s_1 p_1 y_1(z))$?
  - How to prevent sector-2 firms from pretending to be sector-1 firms?
  - Would politicians have the ability to stand up to special interest groups in sector 2?

- Some (e.g. see writings of Anne Krueger, Jagadish Bhagwati,...) have the view that industrial policy shouldn't be tried even if know γ_1/γ_2—that "government failure", in practice, would be worse than *laissez-faire*.

- Others (e.g. Dani Rodrik) have view that industrial policy is no more (or less) subject to this critique than any other area of government policy (e.g. education, infrastructure, monetary policy).

But if we did it, we'd need to know γ_k. How?

- Now imagine a generalization of above model. Multiple countries i, each selling goods in "sub-sector" n (inside industry k). Production within each sub-sector n has the same γ_k.

- This means that supply curve (for sales $S_i^n \equiv y_i^n p_i^n$) for each i and n is:

$$\ln S_i^n = (1 + \epsilon^s) \ln p_i^n + \ln \eta_i^n \quad (1)$$

- ...where ϵ^s is the elasticity of supply in this industry. In the above model this would be given by $\epsilon^s = -1/\gamma_k$

- Suppose also that each i-n pair sells to a bunch of foreign countries j with demand curve (for expenditure D_{ij}^n) given by:

$$\ln D_{ij}^n = \delta_j^n + \delta_i^n + (1 - \epsilon^d) \ln p_i^n + \chi_{ij}^n \quad (2)$$

- Then substituting (1) into (2) we have:

$$\ln D_{ij}^n = \delta_j^n + \delta_i^n + \left(\frac{1 - \epsilon^d}{1 + \epsilon^s} \right) \ln S_i^n + \phi_{ij}^n$$

How to estimate γ_k?

$$\ln D_{ij}^n = \delta_j^n + \delta_i^n + \left(\frac{1-\epsilon^d}{1+\epsilon^s}\right) \ln S_i^n + \phi_{ij}^n$$

- This equation relates bilateral sales to any given destination/buyer (D_{ij}^n) to *total* sales ($S_i^n \equiv \sum_j D_{ij}^n$).
- With an estimate of ϵ^d, we could use this regression to estimate $\epsilon^s = -1/\gamma_k$.
- But doing so would need to IV for S_i^n using a source of the industry's total sales that is not correlated with the error ϕ_{ij}^n, which includes:
  - The demand-shifter χ_{ij}^n for destination j
  - The supply-shifter η_i^n
- What could work? One option is to use some *other* buyer's demand shifter (i.e. for some buyer that is not the j in this equation). And one very important such "other" buyer would be that from *home* demand.

Costinot, Donaldson, Kyle and Williams (QJE 2019)

- CDKW (2019) do exactly this in context of global pharmaceutical industry (c. 2012). Why Pharma?

- Obviously a big tradable industry. And many countries (high-income and middle-income) seem to want to subsidize Pharma development and production.

- Additional (econometric) advantage:
    - Plausibly exogenous shifter of home-country demand coming from demographic composition of home country.
    - Why? Demand for Pharma products is partially determined by demographics, since demographics determine the relative burden of various diseases.
    - Demographic variation in drug *demand* unlikely to be correlated with ways that demographics might matter for *production* of drugs.

- CDKW are basically saying: if countries *export* the drugs that their demographic composition needs most then γ_k (for k=Pharma) must be positive

Predicted Disease Burden (Demand-shifter)

CDKW generate (demographically-) "predicted disease burden" (PDB) as follows:

$$(PDB)_i^n = \sum_{a,g} \left[\text{population}_{iag} \times \left(\frac{\sum_{j \neq i} \text{disease burden}_{jag}^n}{\sum_{j \neq i} \text{population}_{jag}} \right) \right]$$

where:

- n = a "drug category" (for which CDKW can measure D_{ij}^n)
- i = 56 drug-producing countries (with lots of p.c. GDP variation)
- j = major drug-buying countries
- a = age groups (0-14, 15-59, 60+)
- g = gender
- data on disease burden comes from WHO; match disease types to n

(NB: this sort of thing is sometimes called a "shift-share" approach. We saw an example in our discussion of Kovak (2013) in the lecture on trade and inequality.)

IV Estimates

TABLE X

SUPPLY ELASTICITY ESTIMATES

	Log (total sales)	Log (bilateral sales)	
	OLS (1)	OLS (2)	IV (3)
Log (PDB)	1.241		
	(0.110)		
Log (total sales)		0.669	0.764
		(0.052)	(0.116)
p-value for $H_0 : \left(\frac{1-\epsilon^x}{1+\epsilon^s}\right) = 1$			.048**
Adjusted R^2	0.789	0.629	0.627
Observations	18,905	18,905	18,905

$$\ln D_{ij}^n = \delta_j^n + \delta_i^n + \left(\frac{1-\epsilon^d}{1+\epsilon^s}\right) \ln S_i^n + \phi_{ij}^n$$

What does this imply about the size of γ_k?

- Estimate in column (3) implies that $\left(\frac{1-\epsilon^d}{1+\epsilon^s}\right) = 0.764$
- CDKW (2019) also estimate $\epsilon^d = 6.2$ (we won't go into the details here, but it's basically just "standard" demand estimation)
- Put together, estimates therefore imply $\epsilon^s = -7.8$. That implies that $\gamma_k = -1/\epsilon^s = 0.13$ for $k =$ Pharma.

Beyond Pharma

- How does this estimate compare more widely?
- Bartelme, Costinot, Donaldson and Rodriguez-Clare (2025) estimate γ_k separately for each 2-digit manufacturing sector, in a sample of the largest 39 countries around the world (in pooled 1995-2015 cross-sections)
- Instrument for size of each country-industry using idiosyncratic demand differences

BCDR (2025) estimates of γ_k for all of manufacturing

Table 1: Estimates of Sector-Level Scale Elasticities (γ_k)

Sector	OLS (1)	2SLS (2)	Reduced-form (3)	First-stage F-stat (4)	SW F-stat (5)
Food, Beverages and Tobacco	0.19	0.16	0.10	87.20	394.3
	(0.01)	(0.02)	(0.02)		
Textiles	0.14	0.12	0.06	56.70	349.9
	(0.01)	(0.01)	(0.02)		
Wood Products	0.13	0.11	0.05	15.50	210.7
	(0.01)	(0.02)	(0.01)		
Paper Products	0.14	0.11	0.05	55.60	661.9
	(0.01)	(0.02)	(0.01)		
Coke/Petroleum Products	0.09	0.07	0.03	14.20	299.1
	(0.01)	(0.01)	(0.01)		
Chemicals	0.23	0.20	0.17	31.10	335.8
	(0.01)	(0.02)	(0.02)		
Rubber and Plastics	0.29	0.25	0.22	39.13	436.0
	(0.02)	(0.03)	(0.03)		
Mineral Products	0.16	0.13	0.08	40.50	405.0
	(0.01)	(0.02)	(0.01)		
Basic Metals	0.13	0.11	0.07	14.40	254.0
	(0.01)	(0.01)	(0.01)		
Fabricated Metals	0.16	0.13	0.07	57.10	421.1
	(0.01)	(0.02)	(0.01)		
Machinery and Equipment	0.15	0.13	0.07	66.40	401.6
	(0.01)	(0.01)	(0.01)		
Computers and Electronics	0.10	0.09	0.04	18.60	290.5
	(0.01)	(0.01)	(0.01)		
Electrical Machinery, NEC	0.11	0.09	0.03	45.90	419.5
	(0.01)	(0.01)	(0.01)		
Motor Vehicles	0.17	0.15	0.15	39.80	390.2
	(0.01)	(0.01)	(0.02)		
Other Transport Equipment	0.17	0.16	0.11	24.00	381.6
	(0.01)	(0.02)	(0.02)		

Estimates range from $\gamma_k = 0.07$ to $\gamma_k = 0.25$. So span CDKW's Pharma estimate of $\gamma_k = 0.13$.

Back to optimal policy

- Recall, optimal industrial policy subsidy in our 2-industry example depended on γ_1/γ_2
- BCDR (2019) calculate the gains (to each country) of moving from $s_k = 0$ to $s_k = s_k^*$, with extensions to the earlier simple model:
  - Multiple industries: 15 within manufacturing, plus set of non-manufacturing industries
  - Need to know γ_k in the non-manuf industries too! But by setting those $\gamma_k = 0$ will (probably) get upper-bound on effects of industrial policy.
  - More realistic demand than Cobb-Douglas, match each country's data
  - International trade in all goods, with realistic costs of trading
  - Multiple factors of production
  - Input-output linkages (industries use inputs bought from output of other industries/countries)

BCDR (2019): Gains from Optimal Industrial Policy

Table 2: Gains from Optimal Policies, Selected Countries

Country	Optimal Policy (1)	Ind. Policy Only (2)	Trade Policy Only (3)	Gains from Trade Policy (4)	Gains from Ind. Policy (5)
United States	0.55%	0.21%	0.15%	0.34%	0.40%
China	0.66%	0.36%	0.15%	0.30%	0.51%
Germany	0.91%	-0.13%	0.18%	1.04%	0.73%
Ireland	1.56%	-1.49%	0.26%	3.05%	1.31%
Vietnam	1.41%	0.44%	0.62%	0.97%	0.79%
Avg., Unweighted	**1.05%**	**0.10%**	**0.36%**	**0.95%**	**0.69%**
Avg., GDP-weighted	**0.71%**	**0.14%**	**0.20%**	**0.57%**	**0.51%**

Notes: Each column reports the gains, expressed as a share of initial real national income, that could be achieved by each type of policy. See the text for detailed descriptions of the exercises. Column (4) equals the difference between columns (1) and (2); likewise, column (5) is defined as the difference between columns (1) and (3).

Note: these are pretty idealistic scenarios in column (1): planner knows all γ_k's perfectly, and executes s_k^* perfectly; can avoid damage to "terms of trade" by pursuing optimal trade policy too; and, no policy retaliation by foreign countries.

Summary and Next Lecture

- Summary of part I of our coverage of industrial and place-based policy:
  - The logic for IP/PBP is clear: find the industries/places k with the biggest externalities (γ_k) and subsidize (big s_k) them to correct those externalities
  - As with all the productive market failures we've covered in this class, substantial uncertainty remains about the sizes of γ_k
  - But at current estimates of γ_k's, it would be hard to expect large/transformative ("miracle") effects from IP/PBP

- Part II of our coverage of industrial and place-based policy will dig deeper into empirical work on this topic

- Will examine what happened when various countries have tried IP/PBP (or experienced events that might mimic such policies)

- The theory developed in part I will be useful for interpreting what we can learn from such studies in Part II