NBER WORKING PAPER SERIES

ON THE NEGATIVE CONSEQUENCES OF LOW-WAGE OFFSHORING FOR INNOVATION

Wulong Gu
Alla Lileeva
Daniel Trefler

Working Paper 35167
http://www.nber.org/papers/w35167

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2026

We are indebted to Danny Leung, former Director at Statistics Canada, for suggesting the appropriate data and assembling it. We benefited from seminar participants at the Bank of Canada, the Bank of Mexico, Cambridge University the Canadian Economic Association meetings, ERWIT, HKU, LSE, NBER, Statistics Canada, Tsinghua University, UCL, and York University. We benefited from helpful discussions with Costas Arkolakis, Andy Bernard, Kirill Borusyak, Loren Brandt, Paola Conconi, Dave Donaldson, Teresa Fort, Aldo Hernandez, Beryl Jiang, Amit Khandelwal, Kalina Manova, Dalia Marin, Isabelle Mejean, Marc Melitz, Ben Sand, Walter Steingress, Ben Tomlin, Florian Trouvain, and Eric Verhoogen. This research is supported by the Social Sciences and Humanities Research Council of Canada (SSHRC), Grant #435-2016-0185. The views expressed here do not reflect the views of the Bank of Canada, Statistics Canada, or the National Bureau of Economic Research.

On the Negative Consequences of Low-Wage Offshoring for Innovation
Wulong Gu, Alla Lileeva, and Daniel Trefler
NBER Working Paper No. 35167
May 2026
JEL No. F1, F14, O3

ABSTRACT

Conventional wisdom holds that offshoring intermediates to China stimulates innovation. This is not entirely compelling. On the one hand, (a) offshoring lowers marginal costs and expands sales, thereby increasing the returns to innovation, especially for large firms. On the other hand, (b) offshoring low-quality intermediates reduces the costs of older-generation products, thereby reducing the returns to innovating into newer generations. We examine these two opposing forces over 2002-2011 for 6,024 Canadian firms. Our empirical strategy regresses measures of innovation, such as R&D, on imports of intermediate inputs. To address endogeneity, we construct a model-consistent shift-share instrument whose shocks are the often-dramatic improvements in the quality of HS6 Chinese intermediate inputs. We find that greater offshoring reduced R&D spending over 2002-2011 by 15% as (1) firms engaged in R&D in 2002 reduced their expenditures, and (2) firms not initially engaged in R&D were discouraged from starting up new R&D projects. Our model explains these findings: Rising quality of Chinese intermediates is a positive supply shock (rather than a negative China shock) that raises profits for all offshorers, raises innovation for the largest offshorers (channel a above), and lowers innovation for all other offshorers (channel b). These predictions are confirmed in the data.

Wulong Gu
Statistics Canada
wulong.gu@canada.ca

Alla Lileeva
York University
lileeva@yorku.ca

Daniel Trefler
University of Toronto
Rotman School of Management
and NBER
dtrefler@rotman.utoronto.ca

A online appendix is available at http://www.nber.org/data-appendix/w35167

1 Introduction

Starting in the mid-1990s, global supply chains expanded at a rapid pace, with low-wage economies such as China emerging as major suppliers of cheap intermediate inputs to high-wage economies. 15 years into this transformation, concerns surfaced among economists, engineers and business leaders regarding the long-term consequences of offshoring for the competitiveness of high-wage economies. Pisano and Shih (2012), Berger (2013), Fuchs (2014) and others argued that offshoring was raising the costs and lowering the benefits of innovation. Despite a large body of research on the relationship between innovation and the offshoring of intermediate inputs, the extent to which these concerns have merit remains unclear.

In this paper we bring a basic, first-principles approach to the question. Firms build out supply chains in low-wage economies to reduce costs. The resulting lower costs allow firms to sell more output. Since innovation makes each unit of output more profitable, more output means higher returns to innovation. This basic innovation-scale complementarity insight dates back to Schmookler (1954) and was introduced into the offshoring literature by Glass and Saggi (2001). Following Boler, Moxnes and Ullvtveit-Moe (2015), we refer to this as an *innovation-offshoring complementarity*. In this paper we study whether firms in high-wage countries innovate more as a result of importing intermediate inputs from low-wage countries. Our main finding is the opposite: Firm-level importing of intermediates from low-wage countries reduces firm-level innovation. By implication, the concerns expressed 15 years ago are as relevant as ever.

We are hardly the first to examine innovation-offshoring complementarities. Treating productivity gains as evidence of innovation, many firm-level studies find that imports raise productivity e.g., Amiti and Konings (2007) for Indonesia, Kasahara and Rodrigue (2008) for Chile, Topalova and Khandelwal (2011) and De Loecker, Goldberg, Khandelwal and Pavcnik (2016) for India, and Halpern, Koren and Szeidl (2015) for Hungary. In India, tariff cuts on inputs increased product scope (Goldberg, Khandelwal, Pavcnik and Topalova, 2010), further evidence of innovation. In China imports of capital goods directly raised R&D expenditures (Mo, Qiu, Zhang and Dong, 2021).[1] These effects often stem from technology transfer via imports from advanced economies (Bas and Berthou, 2017; Chen, Zhang and Zheng, 2017). In the above studies, imports were typically sourced from countries that were technologically more advanced. *It is therefore unclear whether similar innovation responses occur when firms in technologically advanced, high-wage countries import from less technologically sophisticated, low-wage countries.*

We are aware of only four studies of the innovation-offshoring complementarity for firms in high-wage countries. For Norwegian firms, Boler et al. (2015) find that R&D and imports from OECD countries are complementary, but find no relationship between R&D and imports from low-wage countries. For French firms, Mion and Zhu (2013) show that imports from China slightly reduced innovation as measured by the share of educated workers in total employment. Also for French firms, Aghion, Bergeaud, Lequien, Melitz and Zuber (2024) find that increased importing from China has no statistically significant effect on patenting. In short, these studies provide no

[1] Brandt, Van Biesebroeck, Wang and Zhang (2017) document additional ways low input tariffs helped Chinese firms.

evidence of a strong innovation-importing complementarity for imports from low-wage countries. In these papers it is not clear whether firms are importing intermediate inputs — the subject of this paper — or other goods. The only paper that is clear on this distinction is the Bernard, Fort, Smeets and Warzynski (2024) study of Danish firms. They focus on offshoring of *core activities*, meaning offshoring of produced final goods, and provide clear evidence of an innovation-offshoring complementarity for Danish firms. In contrast, this paper is about the offshoring of intermediate inputs.

Summary of Main Results

We search for an innovation-offshoring complementarity using a balanced panel of 6,024 Canadian firms during 2002–11. The Canadian firm-level data has production data at the HS6 level which we link to 2002–11 HS6 import data. We are therefore able to tell whether an import is (1) a final good that the firm produced or (2) an intermediate input. We focus on the latter.

Using firm-level data and a single long difference, we regress 2002–2011 changes in R&D on 2002–2011 changes in imports of intermediate inputs from low-wage countries (primarily China and Mexico). We employ a novel shift-share instrument for import changes. We apply the Khandelwal (2010) methodology to US import data to compute 2002–2011 quality changes for each country-HS6 product pair. We then aggregate countries into two groups to obtain quality changes for low-wage countries and for high-wage countries. For our shocks we use low-wage relative to high-wage quality changes by HS6 product. We aggregate shocks to the firm level using 2002 firm-level import shares. We treat instrument validity through the Borusyak, Hull and Jaravel (2022, 2025) lens of 'as good as randomly assigned shocks' where Chinese quality upgrading by HS6 are the shocks. We successfully test instrument validity following their many suggestions.

We regress 2002–2011 long changes in R&D expenditures on 2002–2011 long changes in imports of intermediate imports from low-wage countries and instrument the latter with our quality instrument. Our results are stark. *Rising offshoring of intermediates from low-wage countries over 2002–2011 reduced R&D expenditures by 15%.* This results in equal parts from two effects: (1) Firms who did R&D in 2002 reduced their R&D expenditures and (2) firms who did no R&D in 2002 were discouraged from starting up R&D projects in 2011. Projecting out to 2022, R&D fell by 28%. These negative effects came as a surprise to us and led us in search of an overidentified model — meaning one with empirically testable mechanisms — that could explain it.

A Model of Why and When Offshoring Reduces Innovation

It is clear why offshoring might raise innovation. Offshoring lowers marginal costs and increases sales, spreading the benefits of innovation across more units and raising the returns to innovation. This innovation-offshoring complementarity appears in the empirical work of Lileeva and Trefler (2010) and Boler et al. (2015) and will be a feature of both our model and empirics.

Why offshoring might reduce innovation is less clear. Consider the textbook innovation problem. A firm chooses an innovation intensity a to maximize profits. With probability a the firm succeeds in innovation, produces a high-quality, next-generation product, and earns profits π_H. With probability $1 - a$ the firm fails to innovate, produces a low-quality, old-generation product and earns profits π_L. The firm's problem is $\pi(a) = -a^2/2 + a \cdot \pi_H + (1 - a) \cdot \pi_L$ where $a^2/2$ is the total cost of innovation, a is the marginal cost, and $\pi_H - \pi_L$ is the marginal benefit. Thus

$$a^* = \pi_H - \pi_L.$$

All of this is standard and yet this equation lies at the heart of our argument. In particular, suppose that products are produced with two intermediate inputs, a high-quality domestic input m_H and a low-quality Chinese input m_L. Further, the Chinese input has a quality λ_L that is increasing over time so that the quality-adjusted price of m_L is falling. This raises both π_H and π_L. Thus, offshoring can raise or lower innovation, depending on whether $\pi_H - \pi_L$ rises or falls.

When does innovation fall? The answer lies with Shephard's lemma. Consider the share of the Chinese input in total costs. Let s_L be this share for low-quality, old-generation products and let s_H be this share for high-quality, next-generation products. When $s_L > s_H$, improvements in Chinese quality lower old-generation costs by more than next-generation costs. So $\pi_H - \pi_L$ falls and thus so does innovation. We refer to this innovation mechanism as the *relative-cost channel*. Unlike the innovation-offshoring complementarity channel, scale plays no role.

There are two reasons for expecting $s_L > s_H$. First, Kugler and Verhoogen (2012) show that low-quality output is intensive in low-quality inputs. Likewise, Demir, Fieler, Xu and Yang (2024) show that low-quality buyers match with low-quality suppliers. Both of these point to $s_L > s_H$.

Second, numerous engineering case studies have shown that offshoring to low-wage countries can take so much cost out of older-generation technologies that it becomes unattractive to develop next-generation technologies e.g., Fuchs and Kirchain (2007, 2010), Fuchs, Kirchain, and Liu (2011), Fuchs (2014) and Brandt and Wang (2019). This is most likely to happen when offshored intermediates are a large share of costs i.e., when $s_L > s_H$.[2]

In summary improved offshoring opportunities associated with improved Chinese intermediate input quality have two offsetting effects on innovation. The positive innovation-offshoring channel operates through the scale of a firm's operations. In contrast, the negative relative-cost channel operates through relative costs and Shephard's lemma.

Empirically Testable Implications

While our relative-cost channel is a very general feature of many innovation models, we cannot rule out the possibility of other explanations. We therefore develop and test three additional pre-

[2]An example that features in many case studies because of its economic size and importance deals with semiconductors as intermediate inputs. High-quality semiconductors integrate all functions on a single chip and integration requires significant innovation. In contrast, low-quality semiconductors are 'packaged' by linking simpler chips in a process that requires much less innovation and has come to be dominated by Chinese chip makers. As China has taken the cost out of packaged chips, less innovation is done on integrated chips designed for specific manufactured goods. Fuchs and Kirchain (2007, 2010) document this for low-wage offshoring by US firms in the optoelectronics industry, which includes solar panels, fibre optics, camera sensors, screens for smartphones, laptops and TVs. Similar delays in innovating next-generation technologies occurred in data centre equipment, 5G network equipment, and advanced sensors (Fuchs et al., 2011), in automobile bodies (Fuchs, 2014) and in solar panels (Brandt and Wang, 2019).

dictions that help distinguish our channel from alternatives.

First, much of the existing literature on the impact of the China Shock on innovation treats imports as a *demand shock*. Demand shocks reduce profits. See Bloom, Draca and Van Reenen (2015), Autor, Dorn, Hanson, Pisano and Shu (2020) and Yang, Li and Lorenz (2021). In contrast, our intermediate-inputs shock is a *supply shock*. Supply shocks raise profits. Empirically we document rising profits. Thus, our results cannot be explained away by import competition or other explanations involving falling profits.[3]

Second, the model correctly predicts that the positive innovation-offshoring complementarity channel is present in our data, but only where scale is important i.e., only for firms with large innovation step sizes that substantially grow demand. Equating large innovation step sizes with the innovations of multinationals (MNEs) and small step sizes with the innovations of non-MNEs, we find that *all* of our estimated negative impact of low-wage offshoring on innovation is driven by non-MNEs, just as predicted by the model. A caveat to this result is the large standard errors on results for MNEs.

Third, if our negative innovation finding is being driven by the relative-cost channel then the rising quality of Chinese intermediate inputs must be driving firms to substitute away from high-quality domestic inputs and towards Chinese inputs. We find strong evidence of such substitution. This finding also speaks to evolving global supply chains. Our empirics explain most of the 2002–2011 supply-chain restructuring of Canadian firms, meaning their shift away from US suppliers and to Chinese suppliers. There should also be substitution away from domestic labour and towards Chinese inputs, in part because domestic labour is being replaced by the Chinese labour embodied in Chinese intermediate inputs. We find evidence of this: Rising low-wage offshoring reduced firm-level employment by 7%.

Our results are based on IV and are subject to the standard critique of difference-in-difference (DiD) estimation that level effects are not identified or, equivalently, that general-equilibrium effects are ignored. This critique argues for the use of calibration to estimate whether general equilibrium effects dominate DiD results. We show *analytically* that general equilibrium calibration must produce estimates of innovation impacts that are larger (more negative) than our IV/DiD estimates. IV/DiD underestimates the negative impacts on innovation. This analytic result supports our reliance on econometrics.

Related Literature

For trade economists, the most important link between innovation and low-wage offshoring must operate through comparative advantage specialization: China's integration into the world economy pushed rich countries toward innovation-intensive sectors (Arkolakis, Ramondo, Rodríguez-Clare and Yeaple, 2018) and pushed rich-country firms to specialize in their innovation-intensive products (Bernard et al., 2024).

Beyond comparative advantage, many mechanisms link offshoring to innovation through

[3]In section 12, we dig deeper into this by adding seven different controls for changes in product-market competition, including Bartik controls for import competition. This has no effect on our conclusions.

knowledge flows. Offshoring can weaken innovation by separating R&D from production (Naghavi and Ottaviano, 2009; Fort, Keller, Schott, Yeaple and Zolas, 2020; Liu, 2024), but offshoring may also raise the returns to R&D through global knowledge sourcing (Bilir and Morales, 2020). Finally, Branstetter, Chen, Glennon and Zolas (2021) use a natural experiment to show that offshoring via FDI reduces patenting and shifts innovation toward process patents.[4] We will explore some of these mechanisms empirically by looking at the contracting out of R&D, technology licensing, and changes in the number of R&D workers.[5]

The paper is organized as follows. Sections 2–4 describe the model. Section 5 describes the Canadian offshoring and R&D setting. Section 6 describes the data. Sections 7–8 describe the regression model, instruments, and instrument validity. Sections 9–10 describe the negative results for employment and R&D. Section 11 examines mechanisms. Section 12 shows that our results cannot be explained by import competition.

2 A Model

We consider a world with two countries, Canada and China. We assume that Canada is a small open economy in the sense of Melitz and Redding (2014), here meaning that changes in Canada have no impact on Chinese wages w^*. Let w be Canadian wages, let L and L^* be endowments of labour in Canada and China respectively, and let $Y = wL + w^*L^*$ be world income.

Consumers: Preferences over a single differentiated final good are internationally identical and given by

$$U = \left\{ \int_{\omega \in \Omega} \left[\lambda'(\omega)^{\frac{1}{\sigma}} q(\omega)^{\frac{\sigma-1}{\sigma}} \right] d\omega \right\}^{\frac{\sigma}{\sigma-1}} , \tag{1}$$

where $q(\omega)$ is quantity of variety ω, $\lambda'(\omega)$ is its quality, and Ω is the set of available varieties. Since our focus is on traded intermediate inputs, for simplicity we assume that final goods are internationally traded with zero trade costs and no exporting fixed costs so that Ω is the same in both countries. Demand for variety ω is

$$q(\omega) = \lambda'(\omega) p(\omega)^{-\sigma} Y P^{\sigma-1} \tag{2}$$

where $P = \left(\int_{\omega \in \Omega} p(\omega)^{1-\sigma} d\omega \right)^{1/(1-\sigma)}$.

Producers and Output Quality λ: For any variety ω there are potentially two generations, an old-generation and a next-generation which can only be produced if the firm successfully innovates.

[4] See also Chikis, Kleinman and Prato (2025) and Martin and Mejean (2014).

[5] Turning to our negative employment finding, Mion and Zhu (2013) and Aghion et al. (2024) find little evidence that increased importing from China affected employment. Bernard et al. (2024) find that offshoring reduced overall employment. Marin (2010) finds that multinational firms in Austria and Germany in the 1990s outsourced skill-intensive activities to Eastern Europe which led to a fall in skilled-worker wages.

An entering firm pays a sunk cost f_e, draws a baseline productivity φ, and draws a next-generation quality improvement λ that applies if the firm successfully innovates.[6] φ and λ are drawn from a bivariate distribution $G(\varphi, \lambda)$ with continuous first derivatives. After observing (φ, λ), the firm decides whether or not to operate. If it operates it pays a fixed cost f and chooses an R&D intensity $a \geq 0$. With probability $p(a)$, R&D is successful and the firm's quality is λ. Otherwise, quality is 1. In equation (2), $\lambda' = \lambda$ if innovation is successful and $\lambda' = 1$ otherwise.

Let L and H index the Low-quality old generation and the High-quality next generation, respectively. The production function for generation $k = L, H$ is

$$q_k(\varphi) = \varphi \left[\alpha_{lk} \left(l \right)^{\frac{\gamma-1}{\gamma}} \; + \; \alpha_{Hk} \left(m_H \right)^{\frac{\gamma-1}{\gamma}} \; + \; \alpha_{Lk} \left(\lambda_L m_L \right)^{\frac{\gamma-1}{\gamma}} \right]^{\frac{\gamma}{\gamma-1}} \quad \text{for } k \in \{L, H\} \,. \tag{3}$$

There are three inputs: labour l, a high-quality input m_H, and a low-quality input m_L with quality λ_L. $\gamma > 1$ governs the elasticity of substitution.

Intermediate Inputs and their Quality λ_L: We assume that China does not have the capability to produce m_H (as in Sutton and Trefler 2016) and that Canadian suppliers of m_L have sufficiently low productivity that they cannot compete with their Chinese counterparts. In equilibrium m_H is thus only produced in Canada and m_L is only produced in China.

One unit of m_L requires one unit of Chinese labour. m_L is a homogeneous good, is sold on competitive markets, has quality λ_L, and therefore has a quality-adjusted price w^*/λ_L. The quality-adjusted price in Canada is $w^*\tau/\lambda_L$ where τ is an iceberg trade cost.

We assume for simplicity that m_H is produced with one unit of labour and without loss of generality normalize its quality to one. m_H is also a homogeneous good sold on competitive markets and hence has a Canadian quality-adjusted price of w. We assume that the quality adjusted price of m_H is higher than the quality-adjusted price of m_L. That is:

Assumption 1 $w > w^*\tau/\lambda_L$.

The Cost Function: The dual of equation (3) is

$$c_k(w, w^*\tau/\lambda_L) = \varphi^{-1} \left[\alpha_{lk}^{\gamma} \, w^{1-\gamma} \; + \; \alpha_{Hk}^{\gamma} \, w^{1-\gamma} \; + \; \alpha_{Lk}^{\gamma} \left(w^*\tau/\lambda_L \right)^{1-\gamma} \right]^{\frac{1}{1-\gamma}} \quad \text{for } k \in \{L, H\} \,.$$

It will help in what follows to write this in a slightly no-standard way. We choose the normalization $\alpha_{lk}^{\gamma} + \alpha_{Hk}^{\gamma} + \alpha_{Lk}^{\gamma} = 1$. Define $\alpha_k = \alpha_{Lk}^{\gamma}$. Then

$$c_k = \frac{w}{\varphi} \, \Lambda_k^{1/(1-\sigma)} \qquad \text{where} \quad \Lambda_k \equiv \left[(1 - \alpha_k) + \alpha_k \left(\frac{w\lambda_L}{w^*\tau} \right)^{\gamma-1} \right]^{\frac{\sigma-1}{\gamma-1}} . \tag{4}$$

α_k and Λ_k are central in what follows.

[6] In terms of the quality ladders literature (e.g., Grossman and Helpman, 1991), λ is a firm-specific or heterogeneous innovation step size.

Λ_k captures how costs fall as λ_L rises. By Shephard's lemma, the generation-k cost share of the low-quality input is $s_k = \partial \ln c_k / \partial \ln(w^*\tau/\lambda_L)$. By assumption 1, this is increasing in α_k.[7] Holding $w^*\tau$ fixed, Shephard's lemma can equivalently be expressed as

Lemma 1 (Shephard's Lemma)

$$\frac{\partial \Lambda_k}{\partial \lambda_L} = s_k(\sigma - 1)\,\frac{\Lambda_k}{\lambda_L} \qquad \textit{or} \qquad \frac{\partial \ln \Lambda_k}{\partial \ln \lambda_L} = s_k(\sigma - 1)\,.$$

The trivial proof appears in appendix A.1.

Motivated by the empirical work in Kugler and Verhoogen (2012), Demir et al. (2024), Fuchs (2014) and the extended discussion in footnote 2, we assume that low-quality old-generation output is relatively more intensive in the low-quality input:

Assumption 2 $\alpha_L > \alpha_H$.

We tie this back to our earlier comments about the central role of Shephard's lemma by observing that s_k is increasing in α_k so that assumption 2 implies $s_L > s_H$. It likewise implies $\Lambda_L > \Lambda_H$ and $c_L < c_H$.

2.1 Profits and the Innovation Decision

In what follows we treat $w = 1$ as the numeraire. If a firm successfully innovates then it produces a next-generation variety, charges a price $c_H \frac{\sigma}{\sigma-1}$ and earns profits

$$\pi_H(\varphi, \lambda) = \lambda \Lambda_H \varphi^{\sigma-1} B - f - f_M \tag{5}$$

where $B \equiv \sigma^{-\sigma}(\sigma-1)^{\sigma-1} Y P^{\sigma-1}$ and f_M is the fixed cost of importing the Chinese intermediate input. If a firm produces an old-generation variety, it charges a price $c_L \frac{\sigma}{\sigma-1}$ and earns profits

$$\pi_L(\varphi) = \Lambda_L \varphi^{\sigma-1} B - f - f_M\,. \tag{6}$$

Consider the possible orderings of π_H, π_L and zero.

Case 1: When $\pi_H > \pi_L > 0$, the next-generation variety is more profitable than either the old-generation variety or exit. The firm's innovation problem is then

$$\pi(\varphi, \lambda) = \max_{a \ge 0} \left\{ \; -a + p(a)\pi_H(\varphi, \lambda) + \big[1 - p(a)\big]\pi_L(\varphi) \; \right\}. \tag{7}$$

Case 2: When $\pi_H > 0 > \pi_L$, the firm innovates and if unsuccessful it exits. While this may characterize startups, it is not a feature of the firms in our data so that, in the name of simplicity alone, we make an alternative modelling choice. In order to innovate, the firm must engage with

[7] $s_k = \left(\alpha_k p_L^{1-\gamma}\right) / \left[(1-\alpha_k) w^{1-\gamma} + \alpha_k p_L^{1-\gamma}\right]^{\gamma/(\gamma-1)}$ where $p_L \equiv w^*\tau/\lambda_L < w$.

its suppliers and, in so doing, incur the fixed costs f and f_M. Then if the firm successfully innovates it produces the next-generation product and earns variable profits $\lambda\Lambda_H\varphi^{\sigma-1}B > 0$ while if the firm unsuccessfully innovates it produces the old-generation product and earns variable profits $\Lambda_L\varphi^{\sigma-1}B > 0$. The firm thus chooses a to maximize $-a - f - f_M + p(a)\lambda\Lambda_H\varphi^{\sigma-1}B + \big(1 - p(a)\big)\Lambda_L\varphi^{\sigma-1}B$. Conveniently and by design, this is also what the firm maximizes in equation (7).[8]

Case 3: When $\pi_L > \max(\pi_H, 0)$, the old-generation variety is more profitable than either the next-generation variety or exit so the firm does not innovate ($a = 0$) and its profits are $\pi(\varphi, \lambda) = \pi_L(\varphi)$. We assume $p(0) = 0$ so that the firm's problem is also captured by equation (7).

Case 4: When $0 > \max(\pi_H, \pi_L)$ the firm exits.

Collecting these four cases, the firm's problem can be compactly written as equation (7) and the firm exits when $\pi(\varphi, \lambda) < 0$.

If the firm innovates, optimal innovation $a^*(\varphi, \lambda)$ satisfies $\partial\pi/\partial a = 0$ or

$$\frac{1}{p'(a^*)} = \pi_H - \pi_L \,.$$

We assume $p(0) = 0$, $\lim_{a\to\infty} p(a) = 1$, $p' > 0$, $p'' < 0$, and the Inada condition $\lim_{a\searrow 0} p'(a) = \infty$. Let h be the inverse function of $1/p'$, meaning $h(\frac{1}{p'(a)}) = a$. It follows that $h(0) = 0$ and $h' > 0$. See appendix A.2 for the trivial proof. Using h, we can write the previous equation as

$$a^*(\varphi, \lambda) = h\big(\,\pi_H(\varphi, \lambda) - \pi_L(\varphi)\,\big)\,. \tag{8}$$

Since $h' > 0$, a^* is increasing in $\pi_H - \pi_L$. *This is a central result for what follows.* That it is also a common result in the innovation literature speaks to the robustness of the theorems to come.

2.2 Complete Characterization of the Firm's Problem

Figure 1 fully characterizes the firm's optimal choices. The horizontal axis plots $\varphi^{1-\sigma}$, an inverse measure of productivity, so that less-productive firms lie to the right. The vertical axis measures next-generation quality λ. To the bottom right both productivity and quality are low so that $\pi_(\varphi, \lambda)$ is negative and firms exit. To the bottom left firms are productive but quality is low so innovation has low returns and firms choose not to innovate. To the top, quality is high and firms choose to innovate. See appendix A.3 for a proof.

[8]This modelling choice has absolutely no effect on our main theorem (theorem 2) and serves only to slightly simplify the statement of theorem 1.

Figure 1: The Firm's Innovation Decision $a^*(\varphi, \lambda)$

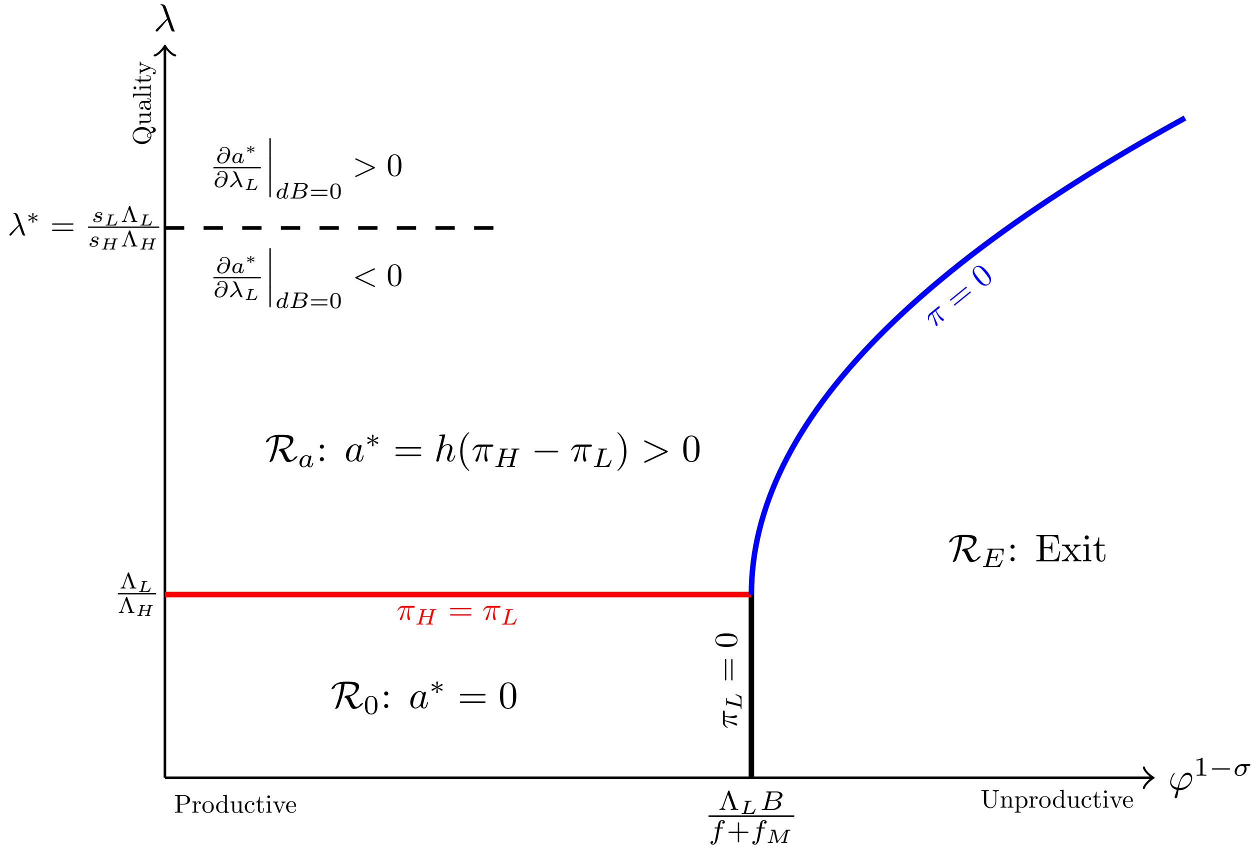

Notes: In the region $\mathcal{R}_E$, firms exit. In the region $\mathcal{R}_0$, firms do not innovate. In the region $\mathcal{R}_a$ firms innovate. Within $\mathcal{R}_a$, a rise in the quality of Chinese intermediate inputs ($\lambda_L \uparrow$) leads to a fall in innovation when $\lambda < \lambda^*$ and a rise in innovation when $\lambda > \lambda^*$. By assumption 2, $\lambda^* \equiv \frac{s_L \Lambda_L}{s_H \Lambda_H} > \frac{\Lambda_L}{\Lambda_H}$ so the region where innovation falls is non-degenerate.

2.3 Comparative Statics without Free Entry

We begin by holding B fixed, which means that we are not allowing for free entry. Consider an increase in the attractiveness of offshoring to China, meaning a rise in the quality of the Chinese intermediate input λ_L. Consider firms for which $a^* > 0$. a^* is increasing in λ_L if $\pi_H - \pi_L = (\lambda \Lambda_H - \Lambda_L)\varphi^{\sigma-1}B$ is increasing in λ_L. Shephard's lemma as expressed in lemma 1 implies

$$\left.\frac{\partial(\pi_H - \pi_L)}{\partial \lambda_L}\right|_{dB=0} = (\lambda s_H \Lambda_H - s_L \Lambda_L)\,\frac{\sigma-1}{\lambda_L}\varphi^{\sigma-1}B\,.$$

It follows that if $a^* > 0$ then

$$\left.\frac{\partial a^*}{\partial \lambda_L}\right|_{dB=0} < 0 \Leftrightarrow \lambda < \lambda^* \equiv \frac{s_L}{s_H} \cdot \frac{\Lambda_L}{\Lambda_H}. \tag{9}$$

The intuition is as follows. For $\lambda < \lambda^*$, firm sales (scale) are small and the innovation-offshoring channel is weak. The relative-cost channel dominates and innovation falls. For $\lambda > \lambda^*$, firm sales (scale) are large and the innovation-offshoring channel is strong so innovation rises.

We summarize this discussion in a formal theorem. Let $\mathcal{R}_a$ ($\mathcal{R}_0$) be the set of firms for which innovation is strictly positive (zero). These sets are easy to define using figure 1. See online appendix D.1 for formal definitions.

Theorem 1 (Innovation without Free Entry)

$$For \;\; (\varphi, \lambda) \in \mathcal{R}_a \;\; \textit{(innovators):} \qquad \left.\frac{\partial a^*(\varphi,\lambda)}{\partial \lambda_L}\right|_{dB=0} < 0 \quad \Longleftrightarrow \quad \lambda < \lambda^* \equiv \frac{s_L}{s_H}\frac{\Lambda_L}{\Lambda_H}\,.$$

$$For \;\; (\varphi, \lambda) \in \mathcal{R}_0 \;\; \textit{(non-innovators):} \quad \left.\frac{\partial a^*(\varphi,\lambda)}{\partial \lambda_L}\right|_{dB=0} = 0$$

Note that a^* is innovation intensity and wa^* is R&D expenditure. With the normalization $w = 1$ the two are interchangeable. Also note that when next-generation products do not use low-quality inputs so that $\alpha_H = 0$, we have $s_H = 0$ and $\lambda^* = \infty$ i.e., *all* firms reduce their innovation.

3 Innovation in General Equilibrium

It is not clear whether the results in theorem 1 hold in general equilibrium, meaning when there is free entry so that B responds to changes in λ_L. We now show that general equilibrium effects via B *strengthen* the conclusion that improved offshoring opportunities can reduce innovation.

3.1 General Equilibrium

We start with a definition of equilibrium. There are markets for (1) Canadian labour, (2) high-quality Canadian intermediate inputs m_H, and (3) varieties of the final good. The prices that clear each of these markets have already been described. See appendix A.4 for a review. We close the model with the free-entry condition:

$$\int_{(\varphi,\lambda)} \pi(\varphi,\lambda)dG(\varphi,\lambda) = f_e\,. \tag{10}$$

This condition pins down B. The normalization $w = 1$ pins down income Y. B and Y together pin down the price index P. This completes the description of equilibrium. Welfare is w/P. [9]

3.2 Innovation in General Equilibrium

Consider a rise in the quality of Chinese intermediate inputs λ_L. General equilibrium considerations introduce one new element, the endogeneity of B. Solving for $dB/d\lambda_L$ involves differentiating equation (10) with respect to λ_L while accounting for how the boundaries in figure 1 change.

[9] There are no externalities in this model so Chinese quality improvements are a "free" cost reduction which raises welfare. Online appendix D.5 discusses the welfare implications of offshoring and innovation in light of Atkeson and Burstein (2010) and our no-externalities assumption.

This involves a lot of work. However, the core insight is very simple and we now describe it. Recall that B depends on P and thus captures all of the general equilibrium feedbacks associated with entry and exit. As λ_L rises, marginal costs fall, profits rise, and firms enter. As is standard with the CES price index, entry reduces P.Hence B falls. The formal proof is in appendix B.

Lemma 2 $dB/d\lambda_L < 0$.

As λ_L rises, in the new equilibrium π_L and π_H cannot both rise for then $\int \pi dG > f_e$, a violation of the free-entry condition. Nor can they both fall for then $\int \pi dG < f_e$. So π_L and π_H must move in opposite directions. Since Λ_L rises by more than Λ_H (Shepherd's lemma and $\alpha_L > \alpha_H$), it follows that π_L rises and π_H falls. Hence, $\pi_H - \pi_L$ falls and with it innovation. This is true *for all firms*.

To state this formally, note that while we have written a^* as a function of (φ, λ), it is also a function of Λ_L, Λ_H and B, all of which depend on λ_L. We therefore write $a^*(\varphi, \lambda; \lambda_L)$. With this notation, we can state our theorem. The proof is in appendix B.

Theorem 2 (Innovation in General Equilibrium) *Consider a rise in the quality λ_L of the Chinese intermediate input. All firms with positive R&D expenditures reduce their R&D and all firms with zero R&D expenditures stay at zero. Formally,*

$$\frac{\partial a^*(\varphi, \lambda; \lambda_L)}{\partial \lambda_L} < 0 \;\; \textit{for} \;\; (\varphi, \lambda) \in \mathcal{R}_a \qquad \textit{and} \qquad \frac{\partial a^*(\varphi, \lambda; \lambda_L)}{\partial \lambda_L} = 0 \;\; \textit{for} \;\; (\varphi, \lambda) \in \mathcal{R}_0 \,. \tag{11}$$

The key difference from the no-free-entry result in theorem 1 is that innovation falls even for firms with $\lambda > \lambda^*$.

4 Three Theoretical Observations Relevant for the Empirics

4.1 Econometrics with and without Free Entry

Our econometrics will essentially compare the R&D growth of firms that were as if randomly exposed to large versus small changes in the quality of their Chinese input suppliers. This means that our empirical results are subject to the usual difference-in-difference caveat that they measure differential effects rather than level effects. The issue is that our treated group – firms that do not import intermediates from China – is still subject to general equilibrium effects operating via B. This type of general equilibrium feedback leads many researchers to augment their econometric work with calibration. The next result shows that we can bound the level effects induced by general equilibrium feedbacks even without calibration.

Corollary 1 *Consider a firm that is innovating. In response to a rise in λ_L, general equilibrium feedback effects push to reduce innovation:*

$$\underbrace{\frac{\partial a^*(\varphi, \lambda; \lambda_L)}{\partial \lambda_L}}_{\textit{Total Effect} < 0} = \underbrace{\left.\frac{\partial a^*(\varphi, \lambda; \lambda_L)}{\partial \lambda_L}\right|_{dB=0}}_{\textit{Direct Effect } \gtrless 0} + \underbrace{\frac{\partial a^*(\varphi, \lambda; \lambda_L)}{\partial B}\frac{\partial B}{\partial \lambda_L}}_{\textit{GE Feedback Effect} < 0} \quad \textit{for} \quad (\varphi, \lambda) \in \mathcal{R}_a \,.$$

The left side is the total impact and is negative by theorem 2. The right side decomposes this into (1) the direct effect which may be negative or positive by theorem 1 and (2) general equilibrium feedback effects. The corollary states that these must be negative. Thus, even if our IV results are subject to the usual difference-in-difference concerns about omitted general equilibrium feedback effects, we know analytically that these effects reduce innovation. *Thus, general equilibrium feedback effects not captured by our IV estimates are always negative i.e., we will underestimate the negative impacts on innovation of offshoring intermediate inputs to China.* The simple proof of the corollary appears in appendix B.3.

4.2 Multinationals Extension

Many of the firms in our sample that innovate are multinational enterprises (MNEs). It is easy to extend the model to allow for these. Motivated by Williamson (1975), we assume that next-generation production requires costly coordination with suppliers (e.g., the cost of knowledge transfer) and this cost can be reduced through vertical integration. To vertically integrate, a firm pays an organizational fixed cost f^V and this reduces next-generation marginal costs of the Chinese input from $\phi^O w^* \tau$ to $\phi^V w^* \tau$ where $\phi^O > \phi^V > 1$.[10] Old-generation Chinese inputs face no coordination costs and so cost $w^* \tau$, as before.

Theorem 3 (MNEs) *For sufficiently large f^V there exists a unique cutoff $\lambda^V(\varphi) > \Lambda_L/\Lambda_H$ such that a firm with $\lambda > \lambda^V(\varphi)$ chooses vertical integration and innovation, a firm with $\lambda^V(\varphi) > \lambda > \Lambda_L/\Lambda_H$ chooses outsourcing and innovation, and a firm with $\Lambda_L/\Lambda_H > \lambda$ chooses outsourcing and no innovation.*

The proof is in online appendix D.2. The theorem states that, on average, MNEs have larger λ than non-MNEs. Essentially, MNEs occupy the upper left of figure 1 where both λ and φ are large. Since the innovation-offshoring complementarity dominates the relative-cost channel for large λ (specifically, $\lambda > \lambda^*$), it follows that the innovation-offshoring complementarity is on average stronger for MNEs than non-MNEs. We will examine this empirically.

4.3 The Extensive Margin of Trade

In our model all firms import Chinese intermediate inputs. Yet two-thirds of firms at the start of our sample did not. It is trivial to extend our model to allow for firms that do not import Chinese inputs. We simply assume that a firm can also produce the old-generation product using only the expensive domestic input, thereby avoiding the fixed costs of importing.[11] Now firms to the lower right of figure 1 neither import nor innovate. See online appendix D.3, especially figure A1. We will use this extension to motivate a fixed-cost instrument that captures the extensive margin of trade. This completes the theory section.

[10]This is as in Helpman, Melitz and Yeaple (2004) with imports replacing FDI and f^V replacing FDI fixed costs.

[11]A producer of the old-generation product faces a standard fixed-cost / marginal-cost technology tradeoff. One technology has high marginal costs (domestic intermediates are expensive) and zero fixed costs of importing. The other technology has low marginal costs (Chinese intermediates are cheap) and positive fixed costs of importing. To capture a Melitz-like pecking order, we assume that only the more productive firms can justify incurring the fixed costs of importing. A necessary and sufficient condition for this is $1 + f_M/f > \Lambda_L$.

Figure 2: The Evolution of Canadian Manufacturing Imports, 1992–2016

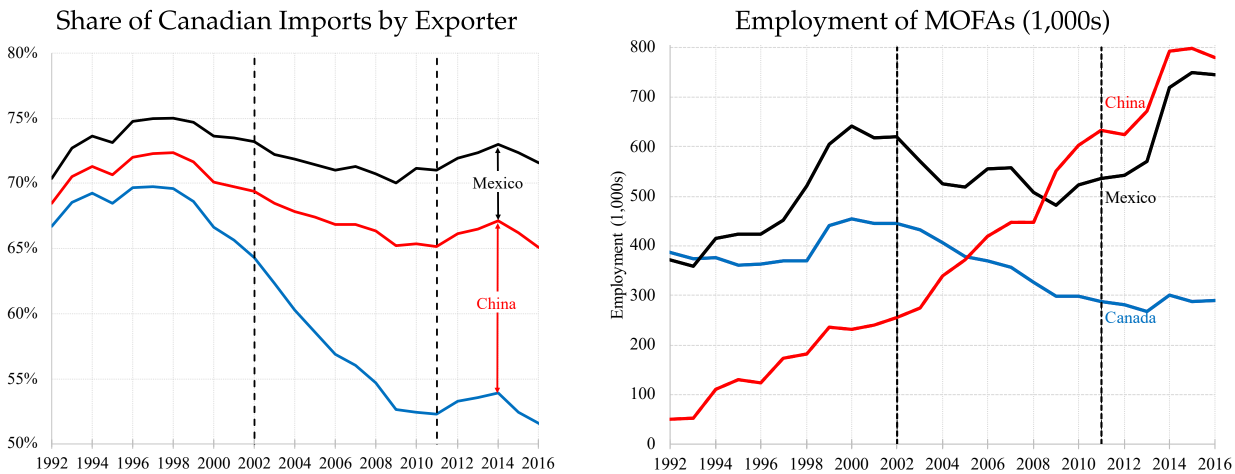

Notes: In the left panel, data are from CANSIM. In the right panel, data are manufacturing employment of majority-owned foreign affiliates of US multinationals (MOFAs) from the Bureau of Economic Analysis "Activities of U.S. Multinational Enterprises".

5 The Canadian Offshoring and R&D Context

By the mid-1990s Canada was heavily integrated into U.S. supply chains. This resulted from the Canada-U.S. Free Trade Agreement and new information technologies that enabled fragmentation e.g., just-in-time. The left panel of figure 2 displays the share of Canadian imports originating from the US (bottom line), the US plus China (middle line), and the US plus China plus Mexico (top line). The panel shows how imports were shifting from high-wage US and towards low-wage Mexico and especially China. This substitution effect is central to our thinking. The two vertical dashed lines demarcate our sample period 2002–2011. The right panel of figure 2 tracks the offshoring decisions of US multinationals, specifically, employment in their manufacturing affiliates operating in Canada, Mexico and China. US multinationals were shifting production out of Canada and into China. 8% of Canadian manufacturing employment was lost as a result of this shift. More generally, Canadian manufacturing employment fell by 22% during 2002–2011, almost exactly the same as in the US. Clearly, similar offshoring forces were operating in both Canada and the US.[12]

Canadian manufacturing R&D data are from Statistics Canada's Research and Development in Canadian Industry (RDCI) and covers all firms with R&D expenditures. The RDCI surveys three types of innovation. The first is expenditures for R&D activities done by the reporting firm for

[12]The importance of intermediates for our results likely reflects Canada's unique position in US supply chains. Many of our firms are either US-owned or heavily integrated into the US manufacturing ecosystem. For example, in 2011 over half of Canadian manufacturing employment was controlled by multinationals, especially US multinationals. These firms famously ship semi-processed goods across the border multiple times. We therefore conjecture that our results would be similar to results using US data.

its own use ('intramural' R&D). This is the most common definition of R&D in academic studies. For Canadian manufacturing during 2002–2011, this R&D declined by 9%. The second expenditure is for R&D the firm contracted out, meaning payments for R&D done by other organizations including parents, affiliates, subsidiaries, and unrelated firms ('extramural' R&D). Contracting out declined by 32%. The third expenditure is payments for the acquisition of knowledge owned by other firms. In our data it is dominated by payments for patents (purchases, licensing) and payments for technical assistance. This expenditure collapsed by 100%. Aggregating all three measures, Canadian manufacturing R&D fell by 23% over 2002–2011. This discussion highlights that we are using a broader measure of R&D than is typical and that Canadian manufacturing R&D fell dramatically over 2002–2011.[13] [14]

6 The Data

We use Canadian firm-level data for the period 2002–2011. Firm-level data come from six data sets maintained by Statistics Canada. (1) Data on employment and other standard variables are from the Annual Survey of Manufactures. (2) Data on imports by HS6 product and country are from the Canadian Imports Registry. (3) Data on R&D expenditures are from the RDCI. (4) Data on profits are from corporate tax returns as reported in the "T2-LEAP" file. (5) Data on whether the firm is foreign-owned are from the *Corporations Returns Act*. (6) Data on whether the Canadian firm has foreign affiliates are from *Activities of Multinational Enterprises in Canada and Abroad*.[15]

We define imported intermediate inputs as follows. Each firm in our sample reports the value of each good it produces. We emphasize that this is production data, not sales data. Thus, for example, goods for resale are excluded. Before 2002, the first six digits of the goods classification correspond to HS6 so we know which HS6 goods the firm did *not* produce pre-2002. If the firm imported one of these non-produced goods in 2002 or later, we refer to it as an imported intermediate input. Restated more formally, for firm f in year t, HS6 good g is an imported intermediate input in year t if the firm did not produce g in either 2000 or 2001, but imported it in year $t \geq 2002$.

There is a source of error in our measure of intermediate inputs. If a firm stops producing a good in year $t \geq 2002$ and starts importing it, we do not include it as an imported intermediate

[13] R&D is an input. The output of R&D is partly captured by a firm's receipts from the sale of its technology. This provides an alternative, output-based lens on innovative activity. Receipts from the sale of technology fell from \$3.4 billion to \$2.6 billion, mirroring what we see in the R&D statistics. (These data are from RDCI annual reports and are not available at the firm level.)

[14] We capture most of a Canadian firm's innovation inputs. What is missing are *some* of the R&D of foreign affiliates of Canadian-owned MNEs. Affiliate innovation is the subject of a number of studies e.g., Bilir and Morales (2020), Branstetter et al. (2021) and Liu (2024). Affiliate innovation is included in our data when the Canadian firm pays its foreign affiliates to conduct R&D or provide technology. If there is no payment, we miss it. To get a sense of magnitudes, for US MNEs in the last decade, about 15% of their global R&D was done by their foreign affiliates, part of this is captured by our data (payments for contracting out and technology), and of what remains, most is likely affiliate R&D aimed at increasing sales in the affiliate's local market (Fan, 2025). As such, we miss very little and what we do miss is reasonably beyond the scope of this paper.

[15] The data in item (6) were first collected in 2008. For our sample, if a Canadian firm had a foreign affiliate in 2008, it had a foreign affiliate in all subsequent years so this is like a fixed firm characteristic. We therefore assume that if a firm in our sample had a foreign affiliate in 2008, it also had one in 2002.

input. Fortunately, this is an advantage for causality: By using pre-2002 production, we purge our measure of endogenous, offshoring-induced changes in product mix. (This logic underpins the use of initial-period weights for Bartik variables.) Thus, our mis-measuring of intermediate inputs has no implications for our causal analysis of the impact of intermediate inputs.[16] [17]

We divide exporters to Canada into two groups depending on their GDP per capita as reported in the World Development Indicators in 2002. Low-wage countries are defined as those with GDP per capita below $5,000. Since the bulk of Canadian imports are from China, Mexico and the US, the only relevant feature of the cutoff is that China and Mexico are below it and the US is above it. See appendix C for the list of rich and poor countries.

We restrict the sample to firms that were present in 2002 and 2011 because the entry and exit dynamics of firms over the Great Financial Crisis are beyond the scope of this paper. Since our main contrast will be between (1) firms that imported from both low- and high-wage countries and (2) firms that imported only from high-wage countries, we restrict the sample to firms that imported in 2002. 99% of these firms imported from a high-wage country in both 2002 and 2011. In contrast, a third of the sample imported from a low-wage country in 2002 and this doubled by 2011. Thus, this paper is not about the effects of importing intermediates, but about the effects of importing intermediates from low- and high-wage countries relative to importing intermediates from high-wage countries only. This keeps the interpretation of the results cleaner. The final sample has 6,024 firms. In 2002, they accounted for about 75% of all manufacturing sales.

7 Regression Model and Instruments

Let f index firms, let t index time and for any variable x_{ft}, define $\Delta x_f \equiv x_{f,2011} - x_{f,2002}$. We are interested in regressions involving a single long change:

$$\Delta y_f = \alpha + \beta \Delta m_{Lf} + \delta X_f + \varepsilon_f \ . \tag{12}$$

y_f will usually be a measure of R&D expenditures. X_f is a vector of initial (2002) firm characteristics described below. Let M_{Lft} be firm f's imports of intermediate inputs from low-wage countries. Since our key mechanism operates via changes in purchases of intermediate inputs, we measure the importance of imported intermediates relative to total material costs $MatCost_{ft}$. We

[16]Previous studies of imports of intermediate inputs classify an HS6 import as an intermediate input based exclusively on its product code e.g., identifying intermediates using the US BEA's "End-Use Category" or the UN's "Broad Economic Category." No firm-level information is used. Thus, a lathe is an intermediate input for all firms, even for a firm producing lathes.

[17]Bernard et al. (2024) use firm-level data to measure 'produced goods'. An HS6 import is a produced good in year t if it was both imported by the firm and produced by the firm in year t. This is a major advance on past research. Also note that their production data allows matching of year t imports with year t production, thus eliminating the type of measurement error discussed above. The share of produced-good imports in total imports is about 5% in both the Danish manufacturing data (Bernard et al. 2024, figure 2) and the Canadian manufacturing data. Online appendix E makes some modest comparisons between the Bernard et al. (2024) data and our data.

Figure 3: The Evolution of Quality, 2002–2011

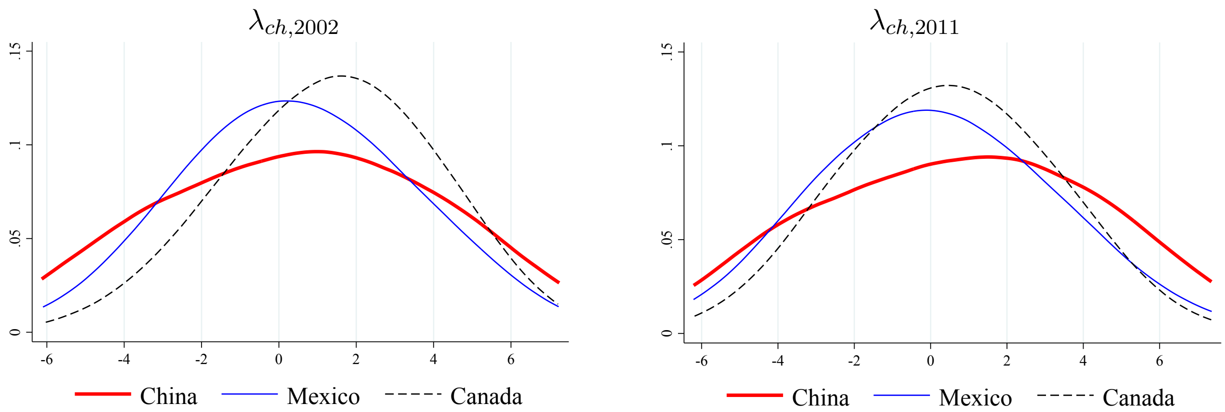

thus define a firm's change in offshoring from low-wage countries by

$$\Delta m_{Lf} = \frac{M_{Lf,2011}}{MatCost_{f,2011}} - \frac{M_{Lf,2002}}{MatCost_{f,2002}} \,. \tag{13}$$

From the theory, this is increasing in λ_L.[18] As shown in online appendix table A7, our results are the same if we instead scale by sales and our results are strengthened if we instead scale by total imports.

7.1 Construction of the Quality and Fixed-Cost Instruments

The model provides two instruments for addressing the endogeneity of offshoring, namely the quality of Chinese intermediates Λ_L and importing fixed costs f_M. We estimate quality λ_L following Khandelwal (2010). We estimate the quality of goods shipped to the US rather than Canada to avoid a potential correlation between goods shipped to Canada and shocks to Canadian firms. We update the Khandelwal HS10 database to 2000–2011 and estimate demands separately for each of 362 NAICS6 industries. The demand estimates are even better than in the original Khandelwal (2010) study. See online appendix F for details.

Let c index countries, let h index HS10 products, let t index time, and denote quality by λ_{cht}. Following Khandelwal, λ_{cht} is measured as the estimated residual (ε_{cht}) from our demand estimation plus the estimated fixed effects λ_{ch} and λ_t. Thus, $\lambda_{cht} \equiv \lambda_{ch} + \lambda_t + \varepsilon_{cht}$.

Figure 3 plots $\lambda_{ch,2002}$ and $\lambda_{ch,2011}$ for c equal to China, Mexico, and Canada. One clearly sees the rightward shift of Mexico and especially China relative to Canada. Between 2002 and 2011, the median value of $\lambda_{Mexico,ht} - \lambda_{Canada,ht}$ rose by 0.5 log points (65%) while the median value of $\lambda_{China,ht} - \lambda_{Canada,ht}$ rose by 1.1 log points (200%). These are massive quality changes.

Within each country-HS10 variety ch, we compute the quality change $\Delta\lambda_{ch} = \lambda_{ch,2011} - \lambda_{ch,2002}$. We aggregate the $\Delta\lambda_{ch}$ across countries to create low- and high-wage country averages and then

[18] $\frac{M_{Lf,2011}}{MatCost_{ft}} = \frac{\alpha_k (p_L)^{1-\gamma}}{\alpha_k (p_L)^{1-\gamma} + \alpha_{Hk}^{\gamma} w}$ is decreasing in $p_L = w^* \tau / \lambda_L$ because $\gamma < 1$ and so increasing in λ_L.

aggregate across goods from HS10 to HS6 so they can be merged with our firm-HS6 data. See online appendix F for details. Let $\Delta\lambda_{Lg}$ and $\Delta\lambda_{Hg}$ be the average quality changes for HS6 good g imported by the US from low- and high-wage countries, respectively. Canada is excluded from the high-wage group. Our shock-level (g) measure of quality change is thus

$$\Delta\lambda_g \equiv \Delta\lambda_{Lg} - \Delta\lambda_{Hg}\,. \tag{14}$$

Differencing across low- and high-wage countries eliminates a final threat to identification.[19]

We can now describe our shift-share instruments. Let s_{fg} be firm f's imports of g in 2002 as a share of the firm's total imports in 2002. $\sum_g s_{fg} = 1$. Our first instrument is

$$\Delta\lambda_f = \sum\nolimits_g s_{fg} \cdot \Delta\lambda_g\,. \tag{15}$$

The model implies that quality satisfies the exclusion restriction. It enters the profit function only via imports of intermediate inputs and thus does not directly affect innovation.

Quality determines how much a firm imports (the intensive margin of trade), but does not fully determine whether a firm imports (the extensive margin). The latter also depends on f_{Mg}, the fixed cost of importing HS6 good g. The smaller is f_{Mg}, the greater is the share of firms that import from low-wage countries. This can be seen in figure 1 by looking at the equations for the boundaries: The smaller is f_{Mg}, the larger is the no-innovation region relative to the innovation region ($\mathcal{R}_0$ relative to $\mathcal{R}_a$). We thus proxy for f_{Mg} using the fraction of all Canadian firms that imported good g from a low-wage country in 2002. Denote this fraction by $p_{Lg,2002}$. Our second shift-share instrument is

$$p_{Lf,2002} = \sum\nolimits_g s_{fg} \cdot p_{Lg,2002}\,. \tag{16}$$

We emphasize that $\Delta\lambda_f$ and $p_{Lf,2002}$ are novel, theory-consistent instruments.

8 Remaining Threats to Identification and Balancing Tests

We have developed model-consistent instruments. Unfortunately, identification may be threatened by features of the environment that stand outside the model. To examine instrument validity we adopt the approaches of Adao, Kolesár and Morales (2019) and Borusyak et al. (2022, 2025) where identification is ensured if the shocks $\Delta\lambda_g$ and $p_{Lg,2002}$ are as good as randomly assigned and display sufficient variability. Appendix D shows that there is abundant shock variability. Here we examine whether the shocks are as-if randomly assigned. There are several threats to this assumption. Positive shocks such as quality improvements to Chinese apparel may be more

[19] We have implicitly adopted the Autor, Dorn and Hanson (2013) identification strategy of using other high-wage countries as a substitute for the target high-wage country (G8 in place of the US in their work, the US in place of Canada here). This is supposed to net out local demand shocks that threaten identification. However, if local shocks are correlated across countries (red shirts are popular this year in every country) then an identification threat remains. The equation (14) differencing across low- and high-wage countries eliminates this threat. Mathematically, differencing the λ_{cht} across country groups, as we do, differences out λ_{ht} components (red shirt popularity). Relatedly, in figure 3 we plot the λ_{cht} after first residualizing them by regressing them on h and t fixed effects.

Table 1: Balancing Tests

Firm Level (f)			Shock/HS6 Level (g)		
	$\Delta\lambda_f$	$p_{Lf,2002}$		$\Delta\lambda_g$	$p_{Lg,2002}$
	(1)	(2)		(3)	(4)
2002 ln(Sales)	-0.08	0.02	2002 ln(Imports)	0.34*	-0.20
	(0.06)	(0.16)		(0.18)	(0.45)
2002 Skill Intensity	0.05	0.12*	2002 ln(Import Price)	0.26	-0.98
	(0.05)	(0.071)		(0.16)	(2.05)
2002 Capital Intensity	-4.94	6.74	2002 ln(Export Price)	0.05	-0.62
	(6.74)	(7.20)		(0.28)	(2.33)
2002 R&D / Sales	-0.00	-0.01	2002 ln(Import Price + Tariff)	0.27	-0.85
	(0.01)	(0.01)		(0.16)	(2.05)

Notes: Each entry is a separate bivariate regression of the variable in the row on the instrument in the column. Columns 1–2 (3–4) are regressions with 6,024 firm-level observations (3,821 HS6-level observations.) Skill intensity is the ratio of non-production workers to production workers. Capital intensity is the ratio of capital to total workers. In columns 1–2, regressions include our five standard controls and NAICS3 fixed effects. In columns 3–4 there are no additional controls. Standard errors are clustered at the NAICS3 level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels.

frequently 'assigned' to firms that are smaller or are less productive or have less-developed global supply chains. To control for this, in our preferred specification, we include 2002 values of

1. the log of employment, a measure of size, and
2. the log of labour productivity,

and three dummies for global supply chain engagement that include whether in 2002 the firm

3. exported,
4. was foreign-owned, and
5. had affiliates outside of Canada.

Collectively, we will refer to these as our 'five standard controls'. We also include 3- or 4-digit NAICS fixed effects.

Table 1 reports balancing tests. Each entry is a separate regression. For the top left entry we regress the log of 2002 sales on our instrument $\Delta\lambda_f$, our five controls, and 3-digit NAICS fixed effects. The regression is at the firm level (6,024 firms or observations). Standard errors clustered at the 3-digit NAICS level appear in parentheses. In column 2 we replace $\Delta\lambda_f$ with our extensive-margin instrument $p_{Lf,2002}$. The instruments are not correlated with firm size in 2002. The remaining rows show that the instruments are uncorrelated with 2002 skill intensity, 2002 capital intensity, and 2002 R&D/sales.

Columns 3–4 repeat the exercise at the shock level, meaning each observation is an HS6 good. The instruments are uncorrelated with the size of the HS6 category as measured by the 2002 log of imports. Nor are the instruments correlated with the 2002 log of import prices, export prices, or delivered import prices (price plus duties paid).[20]

9 Employment

Having discussed instrument validity, we now turn to estimating the causal effects on employment and innovation of offshoring to low-wage countries. Since R&D has both an intensive and extensive margin, we begin with the simpler case of employment. This will help the reader understand our statistical approach.

Consider table 2. There are four columns. In the first there are no covariates. In the second we add our five standard controls. In columns 3–4 we keep the five controls and add either 3- or 4-digit NAICS fixed effects. Reassuringly, estimates are stable across columns in most panels.

Panel A of table 2 is the first stage. It is estimated at the firm level (6,024 observations or firms). The dependent variable is Δm_{Lf} from equation (13) i.e., the 2002–2011 change in firm-level intermediate imports sourced from low-wage countries. The additional independent variables are our two instruments. The first is $\Delta\lambda_f$, the growth in the quality of US imports from low- relative to high-wage countries (equations 14–15). The estimated coefficient across the four columns is positive, indicating that as low-wage quality rises, Canadian firms increase their sourcing of intermediates from low-wage countries. The second instrument is $p_{Lf,2002}$, our inverse measure of fixed costs f_M (equation 16). As expected, firms whose 2002 imports were skewed towards goods with high f_M (low $p_{Lf,2002}$) experienced slower growth in sourcing from low-wage countries. In short, the first stage behaves as predicted by our theory. Panel B is OLS. Again, each observation is a firm. It is unclear what sign to expect. On the one hand, our model implies a reduced-form technology in which, in response to improved low-wage offshoring opportunities, low-quality firms substitute away from domestic labour and towards the foreign labour embodied in imported material inputs. This drives down Canadian employment.[21] On the other hand, firms with trending unobservables (e.g., improving management) may grow both employment and global supply chains, which would bias OLS upwards. Overall, the OLS coefficient is negative but statistically and economically insignificant, consistent with offsetting effects of trending unobservables and our model mechanisms.

Panel C is instrumental variables (what we will call shift-share IV or SSIV) with exposure-robust standard errors corrected as proposed by Borusyak et al. (2022). The exposure-robust weak-

[20] Goldsmith-Pinkham, Sorkin and Swift (2020) offer a different approach to shift-share validity that focusses on the exogeneity of exposure weights s_{fg}. Identification is seen as a difference-in-difference design in which treatments differ because of differences in exposure weights. Goldsmith-Pinkham et al. show that exposure-weight exogeneity is related to parallel trends and can be examined using pre-trend tests. In online appendix G we present the standard pre-trends test (i.e., pooling across shocks) and show that there is no evidence of pre-trends.

[21] In online appendix D.4 we explore the many impacts on employment of rising attractiveness of Chinese offshoring. We show that the dominant forces push for employment declines.

Table 2: Results for the 2002–2011 Log Change in Employment

Fixed effects	No	No	NAICS3	NAICS4
Five controls	No	Yes	Yes	Yes
Panel A. First Stage				
2002-2011 Change in Low-Wage Intermediate $\Delta m_{Lf} = \Delta(M_{Lf} / MatCost_f)$				
Quality Instrument	0.030***	0.030***	0.022**	0.024**
$\Delta\lambda_f$	(0.010)	(0.010)	(0.010)	(0.010)
Fixed Cost Instrument	0.21***	0.21***	0.19***	0.20***
$p_{Lf,2002}$	(0.023)	(0.023)	(0.024)	(0.024)
Panel B. OLS				
2002-2011 Log Change in Employment $\Delta\ln(l_{ft})$				
$\Delta(M_{Lf}/Matcost_f)$	-0.25*	-0.18	-0.08	-0.10
	(0.14)	(0.13)	(0.13)	(0.13)
Panel C. SSIV				
2002-2011 Log Change in Employment $\Delta\ln(l_{ft})$				
$\Delta(M_{Lf}/Matcost_f)$	-3.78***	-2.46***	-2.58***	-2.92***
	(0.62)	(0.54)	(0.72)	(0.74)
Robust F	120	115	78	76
Impact 2002-2011	-0.10	-0.07	-0.07	-0.08
Impact 2002-2022	-0.19	-0.12	-0.13	-0.15
Panel D. SSIV with Shock-Level Demeaning				
2002-2011 Log Change in Employment $\Delta\ln(l_{ft})$				
$\Delta(M_{Lf}/Matcost_f)$	-2.53**	-2.22**	-2.58**	-2.81**
	(1.23)	(1.12)	(1.16)	(1.16)
Cluster Robust F	29	30	29	29
Impact 2002-2011	-0.07	-0.06	-0.07	-0.08
Impact 2002-2022	-0.13	-0.11	-0.13	-0.14

Notes: This table reports estimates of equation (12). Panels A and B are at the firm-level (6,024 firms or observations). Panels C and D are at the HS6 shock-level (3,821 HS6 codes or observations). Panel A is the first-stage regression of $\Delta m_{Lf} = \Delta(M_{Lft}/MatCost_{ft})$ defined in equation (13) on $\Delta\lambda_f$ and $p_{Lf,2002}$ defined in equations (15) and (16). Panels B, C and D report estimates of a regression whose dependent variable is the 2002–2011 log change in employment. Column 1 has no controls. Columns 2–4 include our five standard covariates. Columns 3–4 include 3- and 4-digit NAICS fixed effects. Panel D includes shock-level NAICS3 fixed effects. 'Robust F' is the K-P weak instruments statistic calculated as suggested by Borusyak et al. (2022). 'Impact 2002–2011' is the estimated coefficient times 0.027, the 2002–2011 average over firms of $\Delta(M_{Lf}/MatCosts_f)$. 'Impact 2002–2022' is the estimated coefficient times 0.050, the 2002–2022 average of $\Delta(M_{Lf}/MatCosts_f)$. Robust standard errors are reported in panels A and B. Exposure-robust standard errors are reported in panels C and D. ***, **, and * indicate 1%, 5% and 10% significance levels.

instruments F-statistic is between 76 and 120, well above the Stock-Yogo rule of thumb of 20.[22] Column 3 of panel C is our baseline specification that underlies our headline numbers.

The distribution of shocks may have different means across industries so that we may be exploiting shock variability not just within industries but also across industries. Ideally, we want to exploit within-industry shocks e.g., we want to compare shocks to men's versus women's garments rather than shocks to garments versus autos. We deal with this using the approach in Borusyak et al. (2022), table 4, column 7. That is, we add exposure-weighted sums of NAICS3 fixed effects to force shocks to have the same means across NAICS3 industries.[23] This potentially fundamentally reduces the sample variation in the instruments by eliminating all between-industry variation.[24] Panel D reports the results. The SSIV estimates barely change from what appears in panel C. While the standard errors have grown, we emphasize that panel D has very conservative standard errors by the standards of international trade research: (*a*) We exploit within-NAICS3 shock variation, (*b*) we report exposure-robust NAICS3-clustered standard errors, and (*c*) we eliminate standard-error inflation associated with serial correlation by using a single long difference.

We next turn to the economic magnitudes of our estimates. Our key independent variable is low-wage offshoring as a fraction of total material costs. For the firms in our sample, Δm_{Lf} averaged 0.027 during 2002–2011. Extrapolating to 2022, Δm_{Lf} averaged 0.050 during 2002–2022. These are large changes.[25]

We use the changes 0.027 and 0.050 to conduct two counterfactuals. Consider panel C of table 2. The 'Impact 2002–2011' row is 0.027 $\times\widehat{\beta}^{SSIV}$. The 'Impact 2002–2022' row is 0.050 $\times\widehat{\beta}^{SSIV}$. For our preferred column 3 specification, the rise of low-wage offshoring lowered employment by 7% during 2002–2011 (s.e. = 1.9%) and lowered employment by 13% in 2002–2022 (s.e. = 3.6%).[26]

10 R&D

10.1 The Intensive Margin

Only 20% of our firms have positive R&D in both 2002 and 2011 so we must be cautious about taking log changes of R&D expenditures. We start with the standard, albeit inappropriate, approach

[22]The exposure-robust F-statistics are always larger so we report the empirically more conservative robust F.

[23]Adding exposure-weighted sums of our five controls makes no difference.

[24]This point is clearest if we suppose that each shock-level NAICS3 code has only a single HS6 code. Then shocks are at the NAICS3 level and inclusion of shock-level NAICS3 fixed effects wipes out *all* of the sample variation in the instrument. Shocks then have no identifying power. In practice, there are an average of 182 HS6 codes per NAICS3 cluster and when we include shock-level NAICS3 fixed effects we are identifying solely off the variation of HS6 shocks *within* shock-level NAICS3 clusters.

[25]This may not be obvious because the reader may be unfamiliar with statistics on imports scaled by material cost. To show that these are large numbers, we start by noting that Canadian and US manufacturing are among the most integrated in the world so Canadian imports of intermediates from the US are large by any standard. In 2002, our measure of Canadian intermediates imports from China and Mexico were just one tenth of US levels. By 2011, they were one quarter of US levels, a large surge. This growth is mirrored in figure 2. See appendix C for details.

[26]As discussed above, our model implies a reduced-form technology where, in response to improved low-wage offshoring opportunities, low-quality firms substitute away from domestic labour and towards the foreign labour embodied in imported material inputs. This drives down Canadian employment. We see this both in the theory (online appendix D.6) and in our empirics (online appendix table A5).

Table 3: R&D

Fixed effects	No	No	NAICS3	NAICS4
Five controls	No	Yes	Yes	Yes
	(1)	(2)	(3)	(4)
Δln(1+R&D)				
$\Delta(M_{Lf}/MatCost_f)$	-3.23	-12.67***	-11.66***	-12.74***
	(3.10)	(3.50)	(4.29)	(4.35)
Δ(R&D / Employment): Intensive Margin				
$\Delta(M_{Lf}/MatCost_f)$	-6.36**	-6.23**	-1.43	-0.76
	(2.54)	(2.63)	(3.29)	(3.29)
Impact 2002-2011	-0.15***	-0.11***	-0.08***	-0.08***
	(0.03)	(0.02)	(0.03)	(0.03)
Impact 2002-2022	-0.28***	-0.21***	-0.15**	-0.16***
	(0.05)	(0.05)	(0.06)	(0.06)
Pr(R&D$_{2011}$ > 0): Extensive Margin				
$\Delta(M_{Lf}/MatCost_f)$	-2.48***	-1.12***	-1.07***	-1.15***
	(0.43)	(0.31)	(0.38)	(0.39)
Impact 2002-2011	-0.18	-0.08	-0.08	-0.08
Impact 2002-2022	-0.33	-0.15	-0.14	-0.15

Notes: The dependent variables in the three panels are the 2002–2011 change in the log of R&D, the 2002–2011 change in R&D/employment, and the probability of doing R&D in 2011. The independent variable is the change in low-wage offshoring $\Delta m_L = \Delta(M_{Lf}/MatCost_f)$. It is instrumented by our shift-share instruments $\Delta\lambda_f$ and $p_{Lf,2002}$. The first stage appears in table 2. The weak instruments F is almost identical to those reported in panel 3 of table 2 and exceed 75. In the middle panel the impact is $\frac{d\ln(rd)}{dm_L} \cdot \Delta m_L$ where the derivative is calculated using equation (17) and Δm_L equals either 0.027 or 0.050. In the bottom panel, the impact is divided by the share of firms doing R&D in 2011 (0.380). For example, in column 3, –0.076 = –1.07 · 0.027 / 0.380. Exposure-robust standard errors are reported. ***, **, and * indicate 1%, 5% and 10% significance levels.

of taking the log of one plus R&D. In the first block of results in table 3 the dependent variable is $\ln(1 + rd_{f,2011}) - \ln(1 + rd_{f,2002})$ were rd_{ft} is R&D expenditures by firm f in year t. The first stage is the same as in panel A of table 2 so we only report the SSIV results. While the magnitudes are meaningless (Chen and Roth, 2024), not so the sign. The coefficients are negative. Low-wage offshoring reduces R&D expenditures.

We get at magnitudes using Chen and Roth's suggestion to abandon logs and instead scale by a strictly positive variable. We scale by employment, creating a new dependent variable $\frac{rd_{f,2011}}{l_{f,2011}} - \frac{rd_{f,2002}}{l_{f,2002}}$ where l_{ft} is employment. The results appear in the second block of table 3. A zero coefficient means that low-wage offshoring reduces R&D expenditures by as much as it reduces employment. A negative coefficient means that R&D is impacted by more than employment while a positive coefficient means that R&D is impacted by less than employment (though possibly still negatively). From the second block of results, many of the coefficients are statistically insignificant, implying that low-wage offshoring reduced R&D in proportion to the reduction in employment.

To get more precisely at magnitudes, note that for small changes in R&D and employment it is trivial to show that

$$\frac{d\ln(rd)}{dm_L} \approx \frac{d\ln(l)}{dm_L} + \frac{d\ln(rd/l)}{dm_L} \cdot \frac{l_{2002}}{rd_{2002}} \approx \widehat{\beta}_l + \widehat{\beta}_{rd/l} \cdot \frac{l_{2002}}{rd_{2002}} = -2.58 - 1.43 \cdot 0.28 = -2.98\,. \quad (17)$$

See online appendix D.7. In (17), $\beta_l = -2.58$ is from table 2 (panel C, column 3) and $\beta_{rd/l} = -1.43$ is from table 3 (second block, column 3). When R&D is zero in 2002 this expression blows up. This is just the Chen and Roth problem that for any variable x the percentage change $(x_1 - x_0)/x_0$ is infinite at $x_0 = 0$. Since in this subsection we are interested in intensive-margin responses, we restrict attention to firms with positive R&D in 2002 and evaluate l_{2002}/rd_{2002} at the median of this sample, which is 0.28. Evaluated at this point, the elasticity of R&D expenditures with respect to low-wage offshoring is –2.98 (s.e. = 1.17, p = 0.011). See equation (17).[27] Multiplying –2.98 by Δm_L = 0.027 we obtain the 2002–2011 impact of low-wage offshoring on R&D. From column 3 of table 3, low-wage offshoring lowered R&D by 0.080 log points (s.e. = 0.032, p = 0.01).

We repeat this procedure for each column of table 3 and report the results in the impact rows of the second block. For specifications with NAICS fixed effects, the impacts vary in the tight range of –0.08 to –0.15 log points. Turning to 2002–2022 impacts, we multiply $d\ln(rd)/dm_L$ by Δm_L = 0.050. In our preferred column 3, low-wage offshoring lowered R&D by 0.149 log points (s.e. = 0.058). For specifications with NAICS fixed effects, the 2002–2022 impacts vary between –0.15 and –0.28 log points. *Summarizing based on our preferred specification in column 3, low-wage offshoring reduced R&D expenditures at the intensive margin by 8% in 2002–2011 and by 15% in 2001–2022. These effects are an important conclusion of this paper.*

10.2 The Extensive Margin

The third block of results in table 3 deals with the extensive margin. The dependent variable is a dummy equal to one if the firm did R&D in 2011 and zero otherwise. We interpret this as a linear probability model for the probability that R&D is positive in 2011. The negative coefficients show that low-wage offshoring reduced the probability of doing R&D in 2011. Consider column 3 with its coefficient of –1.07. Low-wage offshoring lowered the probability of doing R&D in 2011 by –1.07 $\times$ 0.027 = –0.029. Since only 38.0% of firms in our sample did R&D in 2011, this is a 7.6% decrease in the number of firms who did R&D in 2011 (0.076 = 0.029/0.380). This number is reported in the bottom block of table 3. The standard error on 7.6% is 2.7%. For 2002–2022, the rise of low-wage offshoring lowered R&D by 14.1% (s.e. = 5.0%). *Summarizing, low-wage offshoring lowered the number of firms doing R&D by 8% during 2002–2011 and by 14% during 2002–2022. These extensive-margin effects are an important conclusion of this paper.*

[27] We compute the standard error by assuming that $\widehat{\beta}_l$ and $\widehat{\beta}_{rd/l}$ are independent.

10.3 Combining Intensive- and Extensive-Margin Responses

Expected R&D equals

$$\mathcal{E}(rd_{ft}) = \mathcal{E}(rd_{ft} \,|\, rd_{ft} > 0) \cdot \mathcal{P}(rd_{ft} > 0)\,.$$

Decomposing this into changes in the intensive margin (changes in R&D conditional on doing R&D) and the extensive margins (changes in the probability of doing R&D) yields

$$\Delta \ln \mathcal{E}(rd_{ft}) = \underbrace{\Delta \ln \mathcal{E}(rd_{ft} \,|\, rd_{ft} > 0)}_{\Delta\text{Intensive Margin}} + \underbrace{\Delta \ln \mathcal{P}(rd_{ft})}_{\Delta\text{Extensive Margin}}\,. \tag{18}$$

If we limit the changes in equation (18) to those that were induced by offshoring intermediates to low-wage countries, the terms in the equation turn out to be closely related to the 'Impact' rows of table 3 and almost identical in magnitude. See online appendix section H for details. *Low-wage offshoring during 2002–2011 reduced R&D by 15%, of which 8% is associated with firms that reduced their R&D expenditures but continued to do R&D and 7% is associated with firms that chose not to do R&D by 2011. For 2002–2022, the corresponding numbers are reductions of 28% (total), 15% (intensive margin) and 13% (extensive margin). These are large and negative impacts of offshoring on innovation.*

10.4 Dynamic Effects on the Transition Path and Steady State

Stepping outside our model, we have also estimated dynamic effects. R&D choices are persistent. A firm that did R&D in 2002 was very likely to do R&D in 2011 and for many years after 2011. As a result, low-wage offshoring that reduced R&D in 2011 likely reduced R&D beyond 2011. We can show that offshoring causally *reduced* the Markov transition probability of starting R&D and thus reduced the share of firms doing R&D both along the transition path and in steady state. The steady-state share fell by a large 11.3%. This result contributes to a small but growing literature showing that the dynamic impacts of offshoring are large e.g., Boehm, Levchenko, Pandalai-Nayar and Toma (2024).

10.5 Details of R&D

The RDCI innovation survey has a lot of detail about innovation. These details appear in table 4. Consider panel A. Column 1 gives the 2002–2011 change in each component of expenditures. Not all components decrease, but many do. Column 2 gives the share of each component in the 2002 total so the reader can focus on the largest components. Column 3 reports results when the dependent variable is the 2002–2011 change in $\ln(1 + y_f)$ for the R&D outcome y_f listed in the row label. Most components were negatively affected, including the largest components. Interestingly, offshoring has a precisely estimated zero effect on R&D contracted out to non-Canadian firms. Thus, our headline negative R&D impacts cannot be explained by the reallocation of R&D from being done within Canadian-based branches of the firm to being contracted out to foreign parents,

Table 4: Detailed Components of R&D Expenditures and Personnel

	2002-2011 Change	2002	y_f=Δln(1+RD$_f$) Coefficient	y_f=Δ(RD$_f$/Employment$_f$) Coefficient	Impact 2002-2011
	(1)	(2)	(3)	(4)	(5)
Panel A. R&D Expenditures Broadly Defined					
1. R&D by reporting firm for itself	**-9%**	**73%**	-11.22***	-0.78	-0.08
			(4.17)	(3.02)	
Wages and salaries	3%	37%	-10.73***	0.46	-0.07
			(4.09)	(2.47)	
Other current expenditures	-19%	32%	-10.52***	-0.97**	-0.08
			(3.26)	(0.40)	
Capital expenditures	-30%	5%	-3.22**	-0.06	-0.07
			(1.58)	(0.05)	
2. R&D contracted out	**-32%**	**15%**	-4.46*	-0.11	-0.07
			(2.55)	(0.19)	
To other Canadian firms	7%	7%	-4.58*	-0.09	-0.07
			(2.55)	(0.19)	
To other non-Canadian firms	-73%	7%	0.04	0.00	-0.07
			(0.09)	(0.00)	
3. Payments for Technology	**-96%**	**12%**	-	-	-
Total 1+2+3	**-23%**	**100%**	-11.85***	-1.33	-0.08
			(4.25)	(3.26)	
Panel B. R&D Personnel					
Total	16%	44,002	-4.98**	-0.01	-0.07
			(2.03)	(0.02)	
Professionals (Scientists + Eng.)	9%	27,305	-7.24***	-0.01	-0.07
			(2.74)	(0.06)	
Technical + admin support staff	9%	16,697	-5.78**	-0.03	-0.07
			(2.47)	(0.03)	
Professionals / Total	-4%	0.62	-0.03	0.02	0.00
			(0.42)	(0.03)	
Average wages	-12%	69,692	1.33	-	0.04
			(1.15)		

Notes: Column 1 is the percentage change in the indicated variable. In panel A, column 2 is the share of the indicated variable in total R&D (1+2+3 sum to 100%). In panel B, column 2 is the number of workers. Column 3 is our standard SSIV regression with independent variable $\Delta(M_{LF}/MatCost_f)$ and dependent variable $\Delta\ln(1+RD_f)$ where RD_f is the indicated variable. Column 4 replaces $\Delta\ln(1+RD_f)$ with $\Delta(RD_f/Employment_f)$. 'Impact 2002–2011' is the impact of offshoring implied by column 4 using equation (17). For 'Professionals/Total', the dependent variable is the change in this ratio (not the log change). 'Average wages' is R&D payroll divided by R&D total personnel. It makes no sense to scale average wages by employment so column 4 is suppressed and column 5 is 1.33×0.027. 'Payments for Technology' is suppressed for confidentiality in columns 3-5. Exposure-robust standard errors are reported. ***, **, and * indicate 1%, 5% and 10% significance levels.

foreign affiliates, foreign subsidiaries, or unrelated firms.[28]

Panel B of table 4 reports results for R&D personnel. The RDCI collects data on the number of professionals (primarily scientists and engineers) and the number of technical and administrative support staff. From column 1, in the raw data R&D personnel increased during 2002–2011. Columns 3–5 show that offshoring led to a 7% reduction in the number of R&D personnel and to no change in the ratio of professional workers to total R&D workers. Average wages (wages and salaries divided by total R&D personnel) rose slightly.

11 Mechanisms

To buttress our claims about the impact of offshoring on innovation, we turn next to documenting in the data the theoretical mechanisms involved. The model highlights three key mechanisms involving (1) profits, (2) substitution towards low-quality inputs, and (3) the opposing forces of the innovation-offshoring complementarity versus the relative-cost channel.

11.1 Mechanism 1: Profits

Improvements in the quality of Chinese intermediate inputs are positive supply shocks. In contrast, in the literature on product-market competition from Chinese imports (e.g., Bloom et al., 2015; Autor et al., 2020), improvements in the quality of Chinese final-good imports are negative demand shocks. This leads to a sharp prediction that distinguishes our model and empirics from previous work on the China shock: *Our positive supply shocks raise profits while China-shock negative demand shocks lower profits.* We should therefore observe that offshoring raises profits. To investigate, we examine the dependent variable

$$\Delta \frac{Profits_f}{Sales_f} = \frac{Profits_{f,2011}}{Sales_{f,2011}} - \frac{Profits_{f,2002}}{Sales_{f,2002}} .$$

We use corporate tax filings to construct profits as revenues less COGS.[29]

The first block of table 5 displays the results. The table has the same format as table 3. There is a statistically significant impact of offshoring on profits. There is also an economically significant impact. In the raw data over the 2002–2011 period, average profitability did not change and its year-to-year standard deviation was 0.9 percentage points. Offshoring raised profits by 0.4 percentage points over 2002–2011, which is about half of a standard deviation. Thus, our results are consistent with positive supply shocks and inconsistent with negative demand shocks (China shocks).[30]

[28]In columns 4–5 the dependent variables are the change in y_f divided by employment and equation (17) is used to calculate the 2002–2001 impacts.

[29]With CES demand, the ratio of gross margins (variable profits) to sales equals $1/\sigma$ and so does not vary across firms. In the data, profit rates vary across firms due to well-known departures from CES (e.g., failure of IIA) and due to firm heterogeneity (e.g., multi-product firms).

[30]There is a reduction in statistical significance because we lose those firms with missing detailed tax data on COGS. As a result, in subsequent tables the profit specifications drop NAICS fixed effects.

Table 5: The Profit and Input-Substitution Mechanisms

Fixed effects	No	No	NAICS3	NAICS4
Five controls	No	Yes	Yes	Yes
	(1)	(2)	(3)	(4)
Δ(Profits$_f$ /Sales$_f$)				
$\Delta(M_{Lf}/MatCost_f)$	0.27***	0.23***	0.16**	0.15**
	(0.06)	(0.06)	(0.08)	(0.08)
Impact 2002-2011	0.007	0.006	0.004	0.004
Impact 2002-2022	0.014	0.012	0.008	0.008
Δ[M$_{Lf}$ / (M$_{Lf}$ + M$_{Hf}$)]: Input Subs.				
$\Delta(M_{Lf}/MatCost_f)$	2.58***	3.91***	3.83***	3.76***
	(0.21)	(0.29)	(0.36)	(0.36)
Impact 2002-2011	0.070	0.106	0.103	0.102
Impact 2002-2022	0.129	0.196	0.192	0.188

Notes: The dependent variables are the 2002–2011 change in the profits-to-sales ratio and the share of imported intermediate inputs originating from low-wage countries. The independent variable is the change in low-wage offshoring $\Delta(M_{Lf}/MatCost_f)$. It is instrumented by our shift-share instruments $\Delta\lambda_f$ and $p_{Lf,2002}$. For the second block of results, the first stage appears in table 2. The weak instruments F is almost identical to those reported in panel C of table 2 and exceed 75. For the first block of results, the number of firms and the number of HS6 shocks is lower, but the first-stage coefficients and weak-instruments F are similar to those in table 2. The impact rows multiply the coefficients by 0.027 (for 2002–2011) and 0.050 (2002–2022). Exposure-robust standard errors are reported. ***, **, and * indicate 1%, 5% and 10% significance levels.

11.2 Mechanism 2: Substitution from High-Quality to Low-Quality Inputs

Improvements in the quality of Chinese intermediates induce substitution from high-quality domestic intermediate inputs to low-quality Chinese intermediate inputs. This is a central model mechanism. We do not observe purchases of domestic intermediates, but given Canada's heavy reliance on US supply chains, we can replace the model's domestic intermediates with intermediates from the US and other high-wage countries. Recall that M_{Lft} and M_{Hft} are imports of intermediate inputs from low- and high-wage countries, respectively. From the quality data used to build our instrument, the quality of M_{Lft} is lower than the quality of M_{Hft} by 0.30 log points in 2002 and by 0.21 log points in 2011. Hence shifts from M_{Lft} to M_{Hft} mirror firm choices to use lower-quality inputs. We therefore expect that increases in λ_L drive a shift in global supply chains from high-wage to low-wage sourcing. To investigate, we estimate our standard regression, but with dependent variable

$$\Delta\frac{M_{Lf}}{M_{Hf}+M_{Lf}} \equiv \frac{M_{Lf,2011}}{M_{Hf,2011}+M_{Lf,2011}} - \frac{M_{Lf,2002}}{M_{Hf,2002}+M_{Lf,2002}}\,.$$

The results appear in the second block of table 5. The coefficients are statistically significant. To get a handle on economic significance, consider the raw data moments for the dependent variable. The move towards sourcing from China and Mexico and away from the US is large in the Canadian data. $\Delta\frac{M_{Lf}}{M_{Hf}+M_{Lf}}$ has an interquartile range of 0.10 and a sample mean of 0.09, indicating a 9–10 percentage point shift in sourcing patterns during 2002–2011. This shift has frequently been noted in the literature on global value chains (e.g., Antràs and Chor, 2022) and was documented in figure 2 for Canada specifically. The 'Impact 2002–2011' line (0.103 in column 3) shows that much of the 9-10 percentage point change in sourcing patterns is attributable to rising Chinese quality.

11.3 Mechanism 3: The Innovation-Offshoring Complementarity versus the Relative-Cost Channel

At the heart of this paper are two propositions. The first is that rising quality of Chinese intermediates is a positive supply shock. Second, this positive supply shock affects R&D in two offsetting ways. (1) Rising λ_L raises R&D through the innovation-offshoring complementarity. This is related to scale. (2) Rising λ_L lowers R&D through the relative-cost channel. This is related to substitution towards low-quality inputs. For large firms, the innovation-offshoring complementarity dominates. For small firms, the relative-cost channel dominates. This tension was expressed mathematically in theorem 1 as

$$\left.\frac{\partial(\pi_H - \pi_L)}{\partial \lambda_L}\right|_{dB=0} > 0 \Leftrightarrow \lambda > \lambda^*\,, \tag{19}$$

an expression that features prominently in the model.[31] We emphasize that this is a unique prediction of the model and that its theoretical generality is apparent from the facts that (1) in many models of innovation $\pi_H - \pi_L$ is the marginal benefit of R&D and so plays a central role and (2) the sign of the above derivative is in large part an implication of a fundamental microeconomic proposition, namely Shepherd's lemma.

To examine the derivative in equation (19) we need to know λ for each firm. While λ can in principle be recovered for a subset of firms using firm revenues conditional on successful innovation, to do so we must observe success and have a selection equation/instrument for success.[32] We therefore appeal to theorem 3, which allows us to recover some aspects of a firm's λ without having to directly observe success.

In theorem 3 we defined a cutoff λ^V above which the firm chooses vertical integration over outsourcing. There are two cases.

[31] Recall that B captures all general equilibrium feedbacks. The derivative in equation 19 (and 20 below) hold B fixed.

[32] Conditional on success, revenues are $r_H(\varphi,\lambda) = \sigma\lambda\Lambda_H\varphi^{\sigma-1}B$, which is the only place where λ enters into the innovation decision. This translates empirically into $\ln(Sales_f) = \ln(\sigma\Lambda_H B) + (\sigma-1)\ln(TFP_f) + \ln(\lambda_f)$ where $\sigma\Lambda_H B$ can be treated as the intercept and $\ln(\lambda_f)$ can be treated as the residual. However, we do not observe success and so do not know what sample of firms to include in this regression. The literature addresses this problem in one of two ways: (1) by measuring success directly using patents (though Canadian firms famously patent very little) and (2) by identifying success stochastically through functional-form assumptions (e.g., $p(a) = \frac{R\&D}{R\&D+1}$) rather than through an additional instrument for success.

1. If $\lambda^V > \lambda^*$ then MNEs all have $\lambda > \lambda^*$ and the probability of $\lambda > \lambda^*$ is greater for MNEs than non-MNEs.
2. If $\lambda^V < \lambda^*$ then non-MNEs all have $\lambda < \lambda^*$ and the probability of $\lambda < \lambda^*$ is greater for non-MNEs than MNEs.

Thus, while MNE status does not perfectly discriminate between $\lambda > \lambda^*$ and $\lambda < \lambda^*$, the model shows that MNE status has discriminatory power. This is also consistent with a large literature on MNEs and innovation.

We begin with profits since these are proximate to the innovation mechanism i.e., the sign of $\partial a^*/\partial \lambda_L$ is determined by the sign of $\partial(\pi_H - \pi_L)/\partial \lambda_L$. From equation (7) and the envelope theorem, if a firm innovates then

$$\frac{\partial \pi}{\partial \lambda_L} = p(a^*)\,\frac{\partial(\pi_H - \pi_L)}{\partial \lambda_L} + \frac{\partial \pi_L}{\partial \lambda_L}\,. \tag{20}$$

$\partial \pi_L/\partial \lambda_L$ is positive and the middle term is positive iff $\lambda > \lambda^*$ (equation 19). This implies that an increase in the quality λ_L of Chinese intermediates raises profits for most non-MNEs and, critically, raises MNE profits by more than non-MNE profits.

To investigate, we stratify the sample into MNEs and non-MNEs. There are 1,018 MNEs, which is 17% of our sample. The first block in table 6 reports the results for profits. Column 1 repeats the results from the specification with our five standard controls (as in tables 2– 5). Columns 2–3 repeat the estimation separately for non-MNEs and MNEs.[33] *The estimated coefficient is much larger for MNEs than non–MNEs, which is central to our mechanism and not obviously predicted by other theories.* The p-value for the difference in coefficients is 0.12, which is not quite significant at the 10% level. In contrast, the difference is economically large as measured by 2002–2011 impacts. Profits rose twice as much form MNEs (0.011) than non-MNEs (0.005). The difference is 0.006 or 0.6 percentage points. Given that over 2002–2011 average profitability did not change and its year-to-year standard deviation was 0.9 percentage points, *0.6 percentage points is a large change.*[34]

Because $\partial a^*/\partial \lambda_L$ inherits the sign of $\partial(\pi_H - \pi_L)/\partial \lambda_L$, we should also see that the impact of λ_L on R&D is larger for MNEs than non-MNEs. In the second block of table 6, the dependent variable is the 2002–2011 change in the ratio of R&D to employment. We can again use equation (17) to recover the impact of offshoring on R&D. These appear in the 'Impact' rows. *We see that R&D fell by 10% for non-MNEs and rose by 1.2% for MNEs. This is an economically major difference.* Although the MNE coefficient is imprecisely estimated, the p-value for the difference of coefficients is 0.07. Thus, subject to the caveat of large standard errors for the MNE results, the difference between the MNE and non-MNE coefficients provides more confirmation that our unique mechanism is at play: The offshoring-innovation complementarity dominates for large-λ firms (MNEs) while the

[33]The first stage is weaker for the MNE sample than for the non-MNE sample. The exposure-robust weak-instruments F is about 12 for MNEs and about 65 for non-MNEs. This accounts for the larger standard errors in the MNE column.

[34]Consistent with our results and suggesting broader validity, there is a small literature showing that the impact of R&D on profits is twice as high for MNEs as for non-MNEs. See Peters and Schmiele (2011) and, for Canada, Gu and Lafrance-Cooke (2020).

Table 6: The Relative-Cost Channel versus the Innovation-Offshoring Complementarity

	All firms	non-MNEs	MNEs	p
	(1)	(2)	(3)	(4)
1. Δ(Profits$_f$ / Sales$_f$)				
$\Delta(M_{Lf}/MatCost_f)$	0.23***	0.18***	0.41***	0.12
	(0.06)	(0.07)	(0.13)	
Impact 2002-2011	0.006	0.005	0.011	
Impact 2002-2022	0.012	0.009	0.021	
2. Δ(R&D / Employment)				
$\Delta(M_{Lf}/MatCost_f)$	-1.43	-2.87	6.17	
	(3.29)	(3.73)	(6.00)	
Impact 2002-2011	-0.08***	-0.10***	0.012	0.07
	(0.03)	(0.04)	(0.051)	
Impact 2002-2022	-0.15***	-0.19***	0.022	0.07
	(0.06)	(0.07)	(0.095)	
3. Δ[M$_{Lf}$ / (M$_{Lf}$ + M$_{Hf}$)]				
$\Delta(M_{Lf}/MatCost_f)$	3.83***	4.10***	2.56***	0.03
	(0.36)	(0.43)	(0.51)	
Impact 2002-2011	0.10	0.11	0.07	
Impact 2002-2022	0.19	0.21	0.13	

Notes: This table repeats our core results in column 1 and re-estimates these results separately for two subsamples: non-MNEs ($N = 5,006$) and MNEs ($N = 1,018$). Each entry is from a separate SSIV regression. Column 4 is the p-value for the test of the null that the MNE and non-MNE coefficients are equal. $p < 0.10$ indicates significance at the 10% level. The test is calculated under the assumption that the two coefficients have zero covariance. In the first and third blocks, the p-values for coefficients and impacts are the same. Exposure-robust standard errors are reported. ***, **, and * indicate 1%, 5% and 10% significance levels.

relative-cost channel dominates for small-λ_L firms (non-MNEs).[35]

We next return to the substitution mechanism. As Chinese quality λ_L rises, firms substitute towards low-quality intermediates from low-wage countries and away from high-quality intermediates from high-wage countries. The third block of table 6 verifies this for both non-MNEs and MNEs. We expect that substitution is greater when low-quality imports are more important, that is, when a firm produces the old-generation product.[36] Thus, we expect the non-MNE coefficient to be larger than the MNE coefficient. This is the case empirically. The 2002–2011 impact

[35]We remind the reader that our standard errors are conservative: they are exposure-robust, we eliminate standard-error inflation associated with serial correlation by using a single long difference, and we include fixed effects in a growth regression.

[36]This follows from assumption 2 ($\alpha_L > \alpha_H$). It holds mathematically as long as imported intermediates are not too large a share of total cost.

is 11 percentage points for non-MNEs and 7 percentage points for MNEs. The difference of 4 percentage points is large ($p = 0.03$). This is again consistent with our unique mechanism.

Concluding this section, we showed that Chinese quality improvements (increases in λ_L) are a positive supply shock, not a negative demand shock. We also documented the presence of input substitution towards low-quality inputs, a key underlying mechanism of the model. Finally, we showed that the unique mechanism of our model, which operates via the sign of $\partial(\pi_H - \pi_L)/\partial\lambda_L$, can be examined by dividing the sample into MNE and non-MNE subsamples and examining their differential responses to increases in λ_L. The differential responses for profits, R&D, and input substitution are all as predicted by the model.

12 Product-Market Competition

It is possible that low-wage offshoring is correlated with product-market competition, in which case our results may be driven more by product-market competition than offshoring of intermediate inputs. We argued that this is unlikely because product-market competition is a negative demand shock that lowers profits.

To further investigate, we introduce three measures of domestic product-market competition and three measures of import competition. Starting with domestic competition, we have NAICS6 data on 4-firm concentration ratios, Herfindahl indexes, and log numbers of firms. We link these to our firms using each firm's NAICS6 major line of business. In column 1 of table 7 we repeat the most important results of our paper. That is, we repeat column 3 of table 3. In columns 2–4 we add in the domestic competition variables. Looking across these columns, our baseline results do not change when we control for domestic competition.

Turning to import competition, there are two possible approaches. We can again use a NAICS6 link between our firms and import competition. However, a finer approach exploits import competition at the HS6 level by creating a shift-share measure of import competition. Let IC_{gt} be a measure of import competition in the product market for HS6 good g. For 2000–2001, we have each firm's output of produced goods. Let s^P_{fg} be the share of HS6 good g in firm f's output. This is the firm's product-market exposure to imports of g. We measure import competition by

$$\Delta IC_f = \sum\nolimits_g s^P_{fg}\left(IC_{g,2011} - IC_{g,2002}\right) .$$

In column 5 of table 7 we measure IC_{gt} as the Canadian tariff against low-wage countries. In column 6 we measure IC_{gt} as Canadian log imports of g from low-wage countries. This may be correlated with observable or unobservable features of the firm and so in column 7 we replace Canadian imports with US imports. This can be interpreted as a reduced-form version of the Autor et al. (2013) approach in that we instrument imports with imports of another rich country.

Looking across columns 5–7 in the top and bottom panels of table 7, it is clear that our results are not driven by product-market competition. In the middle panel the dependent variable is R&D/Employment and, as we discussed, what matters most is not the coefficient, but the equation

Table 7: SSIV Estimates of Coefficients on Δm_{Lf} with Controls for Product-Market Competition

	Baseline	Domestic Competition			Foreign Competition		
	Baseline	CONC4	Herfindahl	log number of firms	Tariffs	Cdn. Low-Wage Imports	US Low-Wage Imports
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δln(1+R&D)							
$\Delta(M_{Lf}/MatCost_f)$	-11.66***	-11.54***	-11.76***	-11.76***	-10.97**	-10.88**	-10.81**
	(4.29)	(4.14)	(4.15)	(4.14)	(4.31)	(4.32)	(4.31)
Δ(Competition)		-1.70**	-1.67*	0.23	-0.71	0.03	0.13
		(0.69)	(0.88)	(0.44)	(3.39)	(0.05)	(0.09)
Δ(R&D/Employment): Intensive Margin							
$\Delta(M_{Lf}/MatCost_f)$	-1.43	-1.36	-1.47	-1.31	-0.13	-0.29	-0.08
	(3.29)	(2.91)	(2.91)	(2.89)	(3.01)	(3.03)	(3.02)
Δ(Competition)	-0.08	-0.69	-1.28**	-0.39	1.11	-0.05	0.03
		(0.51)	(0.64)	(0.32)	(2.29)	(0.04)	(0.07)
Impact 2002-2011	-0.080	-0.08	-0.08	-0.08	-0.08	-0.08	-0.07
Impact 2002-2022	-0.148	-0.15	-0.15	-0.15	-0.14	-0.14	-0.14
$Pr(R\&D_{2011} > 0)$: Extensive Margin							
$\Delta(M_{Lf}/MatCost_f)$	-1.07***	-1.06***	-1.08***	-1.08***	-1.02***	-1.00***	-1.00***
	(0.38)	(0.37)	(0.37)	(0.37)	(0.39)	(0.39)	(0.39)
Δ(Competition)		-0.15**	-0.14*	0.02	-0.12	0.00	0.01
		(0.06)	(0.08)	(0.04)	(0.30)	(0.01)	(0.01)
Impact 2002-2011	-0.076	-0.08	-0.08	-0.08	-0.07	-0.07	-0.07
Impact 2002-2022	-0.141	-0.14	-0.14	-0.14	-0.13	-0.13	-0.13

Notes: Column 1 carries over results from the previous three tables. The remaining columns add different measures of product-market competition. All specifications include our five standard controls and the 2002 value of the dependent variable (to control for both heterogeneity and mean reversion). Robust standard errors (not exposure-robust) are reported in parentheses. ***, **, and * indicated significance at the 1%, 5% and 10% levels using these standard errors.

(17) estimated impact. This appears in the 2002–2011 impact row. In our baseline, offshoring reduced R&D by 8%. In rows 5–7 this number is almost the same. We conclude that our results cannot be explained away by changes in product-market competition from either domestic or foreign competitors.[37]

[37]This paper is too long to properly investigate the impacts of domestic and low-wage competition on R&D. Autor et al. (2020) find negative impacts of Chinese competition. We find negative impacts for domestic competition. We also find some negative impacts for low-wage competition, but only on the intensive margin — the near-zero intensive-margin coefficient in the middle panel shows that R&D fell at the same rate as employment. This negative finding is mirrored in Autor et al.. Bloom et al. (2015) find positive impacts of Chinese competition, though for patenting they cannot rule out zero effects (Campbell and Mau, 2021). Yang et al. (2021) find negative impacts on self-reported process innovation and positive impacts on self-reported product innovation so it is possible that the net effect is small, consistent with our findings.

13 Additional Robustness

Our core measure of offshoring, m_{Lft} in equation (13), scaled low-wage offshoring by material costs. Our estimated impacts are the same when we instead scale by the firm's sales. They are about 50% larger when we instead scale by the firm's total imports. See online appendix table A7.

We do not expect that offshoring intermediate inputs from high-wage countries will have negative effects. Online appendix table A7 shows that high-wage offshoring is always statistically insignificant.[38]

It is possible that our results have something to do with the fragmentation of manufacturing (Bernard and Fort, 2015). The Annual Survey of Manufactures distinguishes between sales associated with manufactured goods versus non-manufactured goods such as services or goods for resale. We compute the 2002–2011 change in the ratio of non-manufacturing sales to total sales and regress it on Δm_{Lft}. From online appendix table A6, the relationship is both economically and statistically tiny. Also, while our measure of intermediate inputs intentionally excludes goods for resale, we have goods for resale at the firm level and, in online appendix table A6, we regress the 2002–2011 log change in goods for resale on Δm_{Lft}. There is no impact of offshoring intermediate inputs on goods for resale.

14 Conclusions

We examined the implications for innovation of offshoring intermediates to China and other low-wage economies. We have 2002–2011 data on 6,024 Canadian firms. Using HS6-level data on both imports and outputs, we constructed a firm-specific measure of imports of intermediate inputs. We regressed a firm's long change in R&D on its long change in imports of intermediate inputs from low-wage economies. We instrumented for imports using novel and model-consistent shift-share instruments. (We tested extensively for instrument validity.) The IV estimates imply that rising offshoring of intermediates from low-wage countries over 2002–2011 reduced R&D expenditures by 15% as (1) firms with positive R&D in 2002 reduced their R&D expenditures and (2) firms with zero R&D in 2002 were discouraged from starting new R&D projects.

In the model, rising quality of Chinese intermediates (1) is a positive supply shock that raises profits, (2) leads to substitution towards low-quality inputs, and (3) affects R&D via an offshoring-innovation complementarity and offsetting relative-cost channel. We provided evidence that these unique mechanisms empirically underpin our negative R&D results.

[38] Letting M_{Hft} be firm f's imports of intermediate inputs from high-wage countries, we measure high-wage offshoring as $\Delta m_{Hf} = \frac{M_{Hf,2011}}{MatCost_{f,2011}} - \frac{M_{Hf,2002}}{MatCost_{f,2002}}$.

References

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Aghion, Philippe, Antonin Bergeaud, Matthieu Lequien, Marc J. Melitz, and Thomas Zuber, "Opposing Firm-Level Responses to the China Shock: Output Competition versus Input Supply," *American Economic Journal: Economic Policy*, May 2024, *16* (2), 249–269.

Amiti, Mary and Joseph Konings, "Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia," *American Economic Review*, December 2007, *97* (5), 1611–1638.

Antràs, Pol and Davin Chor, "Global Value Chains," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. 5, Elsevier, 2022, pp. 297–376.

Arkolakis, Costas, Natalia Ramondo, Andres Rodríguez-Clare, and Stephen Yeaple, "Innovation and Production in the Global Economy," *American Economic Review*, August 2018, *108* (8), 2128–2173.

Atkeson, Andrew and Ariel Burstein, "Innovation, Firm Dynamics, and International Trade," *Journal of Political Economy*, June 2010, *118* (3), 433–484.

Autor, David, David Dorn, and Gordon Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–68.

_ , _ , Gordon H. Hanson, Gary Pisano, and Pian Shu, "Foreign Competition and Domestic Innovation: Evidence from U.S. Patents," *American Economic Review: Insights*, September 2020, *2* (3), 357–374.

Bas, Maria and Antoine Berthou, "Does Input-Trade Liberalization Affect Firms' Foreign Technology Choice?," *The World Bank Economic Review*, 2017, *31* (2), 351–384.

Berger, Suzanne, *Making in America: From Innovation to Market*, Cambridge, MA: MIT Press, 2013.

Bernard, Andrew B. and Teresa C. Fort, "Factoryless Goods Producing Firms," *American Economic Review Papers & Proceedings*, May 2015, *105* (5), 518–23.

_ , _ , Valerie Smeets, and Frederic Warzynski, "Heterogeneous Globalization: Offshoring and Reorganization," WP No. 26854, National Bureau of Economic Research, November 2024.

Bilir, L. Kamran and Eduardo Morales, "Innovation in the Global Firm," *Journal of Political Economy*, 2020, *128* (4), 1566–1625.

Bloom, Nicholas, Mirko Draca, and John Van Reenen, "Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity," *Review of Economic Studies*, 2015.

Boehm, Christoph, Andrei A. Levchenko, Nitya Pandalai-Nayar, and Hiroshi Toma, "Dynamic Models, New Gains from Trade?," Working Paper No. 32565, National Bureau of Economic Research, June 2024.

Boler, Esther Ann, Andreas Moxnes, and Karen Helene Ullvtveit-Moe, "R&D, International Sourcing and the Joint Impact on Firm Perfomance," *American Economic Review*, December 2015, *105* (12), 3704–3739.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

_ , _ , and _ , "A Practical Guide to Shift-Share Instruments," *Journal of Economic Perspectives*, 2025, *39* (1), 181–204.

Brandt, Loren and Luhang Wang, "China's Development of Wind and Solar Power," in Loren Brandt and Thomas G. Rawski, eds., *Policy, Regulation and Innovation in China's Electricity and Telecom Industries*, Cambridge, UK: Cambridge University Press, 2019, pp. 373–418.

_ , Johannes Van Biesebroeck, Luhang Wang, and Yifan Zhang, "WTO Accession and Performance of Chinese Manufacturing Firms," *American Economic Review*, September 2017, *107* (9), 2784–2820.

Branstetter, Lee G., Jong-Rong Chen, Britta Glennon, and Nikolas Zolas, "Does Offshoring Production Reduce Innovation: Firm-Level Evidence from Taiwan," Working Paper No. 29117, National Bureau of Economic Research, August 2021.

Campbell, Douglas L and Karsten Mau, "On "Trade Induced Technical Change: The Impact of Chinese Imports on Innovation, IT, and Productivity"," *Review of Economic Studies*, October 2021, *88* (5), 2555–2559.

Chen, Jiafeng and Jonathan Roth, "Logs with Zeros?: Some Problems and Solutions," *Quarterly Journal of Economics*, May 2024, *139* (2), 891–936.

Chen, Zhiyuan, Jie Zhang, and Wenping Zheng, "Import and innovation: Evidence from Chinese firms," *European Economic Review*, 2017, *94* (C), 205–220.

Chikis, Craig A., Benny Kleinman, and Marta Prato, "The Geography of Innovative Firms," Working Paper No. 34010, National Bureau of Economic Research, July 2025.

De Loecker, Jan, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik, "Prices, Markups, and Trade Reform," *Econometrica*, March 2016, *84*, 445–510.

Demir, Banu, Ana Cecília Fieler, Daniel Yi Xu, and Kelly Kaili Yang, "O-Ring Production Networks," *Journal of Political Economy*, 2024, *132* (1), 200–247.

Fan, Jingting, "Talent, Geography, and Offshore R&D," *Review of Economic Studies*, 2025, *92* (2), 1022–1060.

Fort, Teresa C., Wolfgang Keller, Peter K. Schott, Stephen Yeaple, and Nikolas Zolas, "Colocation of Production and Innovation: Evidence from the United States," Working Paper 2020.

Fuchs, Erica R.H., "Global Manufacturing and the Future of Technology," *Science*, August 2014, *345* (6196), 519–520.

— **and Randolph E. Kirchain**, "Changing Paths: The Impact of Manufacturing Offshore on Technology Development Incentives in the Optoelectronics Industry," Working Paper No. 16, Industry Studies Association, June 2007.

— **and** — , "Design for Location? The Impact of Manufacturing Offshore on Technology Competitiveness in the Optoelectronics Industry," *Management Science*, 2010, *56* (12), 2323–2349.

— **,** — **, , and Shan Liu**, "The Future of Silicon Photonics: Not So Fast? Insights From 100G Ethernet LAN Transceivers," *Journal of Lightwave Technology*, August 2011, *29* (15), 2319–2326.

Glass, Amy Jocelyn and Kamal Saggi, "Innovation and wage effects of international outsourcing," *European Economic Review*, January 2001, *45* (1), 67–86.

Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova, "Imported Intermediate Inputs and Domestic Product Growth: Evidence from India," *Quarterly Journal of Economics*, November 2010, *125* (4), 1727–1767.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, August 2020, *110* (8), 2586–2624.

Grossman, Gene M. and Elhanan Helpman, "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, 1991, *58* (1), 43–61.

Gu, Wulong and Amélie Lafrance-Cooke, "Why Are Multinationals More Productive than Non-multinationals? Evidence from Canada," Analytical Studies Branch Research Paper Series 11F0019M, no. 447, Statistics Canada 2020.

Halpern, László, Miklós Koren, and Adam Szeidl, "Imported Inputs and Productivity," *American Economic Review*, December 2015, *105* (12), 3660–3703.

Helpman, Elhanan, Marc J. Melitz, and Stephen R. Yeaple, "Export versus FDI with Heterogeneous Firms," *American Economic Review*, 2004, *94* (1), 300–316.

Kasahara, Hiroyuki and Joel Rodrigue, "Does the Use of Imported Intermediates Increase Productivity? Plant-level Evidence," *Journal of Development Economics*, August 2008, *87* (1), 106–118.

Khandelwal, Amit, "The Long and Short (of) Quality Ladders," *Review of Economic Studies*, October 2010, *77* (4), 1450–1476.

Kugler, Maurice and Eric Verhoogen, "Prices, Plant Size, and Product Quality," *Review of Economic Studies*, January 2012, *79* (1), 307–339.

Lileeva, Alla and Daniel Trefler, "Improved Access to Foreign Markets Raises Plant-Level Productivity ... for Some Plants," *Quarterly Journal of Economics*, August 2010, *125* (3), 1051–1100.

Liu, Jin, "Multinational Production and Innovation in Tandem," Working Paper No. 24–64, Center for Economic Studies, October 2024.

Marin, Dalia, "The opening up of Eastern Europe at 20 – jobs, skills, and 'reverse maquiladoras' in Austria and Germany," *Bruegel, Working Papers*, 01 2010.

Martin, Julien and Isabelle Mejean, "Low-wage country competition and the quality content of high-wage country exports," *Journal of International Economics*, None 2014, *93* (1), 140–152.

Melitz, Marc J. and Stephen J. Redding, "Heterogeneous Firms and Trade," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Trade*, Amsterdam: Elsevier, 2014, pp. 1–54.

Mion, Giordano and Linke Zhu, "Import competition from and offshoring to China: A curse or blessing for firms?," *Journal of International Economics*, 2013, *89* (1), 202–215.

Mo, Jiawei, Larry D. Qiu, Hongsong Zhang, and Xiaoyu Dong, "What you import matters for productivity growth: Experience from Chinese manufacturing firms," *Journal of Development Economics*, 2021, *152* (C).

Naghavi, Alireza and Gianmarco Ottaviano, "Offshoring and Product Innovation," *Economic Theory*, March 2009, *38* (3), 517–532.

Peters, Bettina and Anja Schmiele, "The Contribution of International R&D to Firm Profitability," ZEW Discussion Papers 11–002, ZEW - Leibniz Centre for European Economic Research 2011.

Pisano, Gary P. and Willy Shih, *Producing Prosperity: Why America Needs a Manufacturing Renaissance*, Boston: MA: Harvard Business Review Press, 2012.

Schmookler, Jacob, "The Level of Inventive Activity," *Review of Economics and Statistics*, May 1954, *36* (2), 183–190.

Sutton, John and Daniel Trefler, "Deductions from the Export Basket: Capabilities, Wealth and Trade," *Journal of Political Economy*, June 2016, *124* (3), 826–878.

Topalova, Petia and Amit Khandelwal, "Trade Liberalization and Firm Productivity: The Case of India," *Review of Economics and Statistics*, August 2011, *93* (3), 995–1009.

Williamson, Oliver E., *Markets and Hierarchies: Analysis and Antitrust Implications*, New York: Free Press, 1975.

Yang, Mu-Jeung, Nicholas Li, and Kueng Lorenz, "The Impact of Emerging Market Competition on Innovation and Business Strategy: Evidence from Canada," *Journal of Economic Behaviour & Organization*, 2021, *181* (C), 117–134.

Appendix

A Model Proofs

A.1 Proof of Lemma 1

Denoting the quality-adjusted price of m_L as $p_L = w^*\tau/\lambda_L$, Shephard's lemma states that $s_k = \partial \ln c_k/\partial \ln p_L$ or, holding w^* and τ constant, $s_k = -\partial \ln c_k/\partial \ln \lambda_L$. Rewrite equation (4) as $\Lambda_k = (\varphi c_k/w)^{1-\sigma}$. Then $\frac{\partial \ln \Lambda_k}{\partial \ln \lambda_L} = (1-\sigma)\frac{\partial \ln c_k}{\partial \ln \lambda_L} = (\sigma-1)s_k$ and $\frac{\partial \Lambda_k}{\partial \lambda_L} = (\sigma-1)s_k\frac{\Lambda_k}{\lambda_L}$.

A.2 Proof that $h(0) = 0$ and $h' > 0$

h is implicitly defined by $h(1/p'(a)) = a$. Hence $\lim_{a\searrow 0} h(1/p'(a)) = \lim_{a\searrow 0} a$ or $h(1/\infty) = 0$ or $h(0) = 0$. $h' > 0$ follows from differentiating $h(1/p'(a)) = a$ to obtain $h'(1/p'(a))[-p''/(p')^2] = 1$ or $h' = -(p')^2/p'' > 0$.

A.3 Appendix Proof of Figure 1

From equation (6), $\pi_L(\varphi) = 0$ can be written as $\varphi^{1-\sigma} = (\Lambda_L B)/(f + f_M)$. This is the vertical line in figure 1. To the left of this line, the firm innovates if and only if $\pi_H > \pi_L$. By equations (5)–(6), the equation $\pi_H(\varphi, \lambda) = \pi_L(\varphi)$ can be written as $\lambda = \Lambda_L/\Lambda_H$, the horizontal line in figure 1. Above it the firm innovates and below it the firm does not. Next consider the region to the right of the vertical line $\pi_L = 0$. If $\pi(\varphi, \lambda) < 0$ the firm prefers exit. If $\pi(\varphi, \lambda) > 0$ the firm prefers innovating. The shape of the $\pi(\varphi, \lambda) = 0$ curve is as follows. At the figure 1 intersection of the $\pi_L = 0$ and $\pi_H = \pi_L$ lines we, must have $\pi_H = 0$. But $\pi_L = 0$ and $\pi_H = 0$ imply $\pi = 0$ i.e., all three figure 1 lines intersect at the same point. Further, $\pi(\varphi, \lambda) = 0$ is upward sloping because the lower is productivity, the higher must λ be in order to maintain profits at zero.

A.4 Market Clearing

Consider the determination of prices in markets for (1) Canadian labour, (2) high-quality Canadian intermediate inputs m_H, and (3) varieties of the final good. For the labour market, $w = 1$. For the high-quality intermediates market, we assumed m_H is produced only with labour, is subject to constant returns to scale, and is sold in perfectly competitive markets so its price is 1. In markets for final-good varieties, price equals a constant markup over marginal cost. Since $w = 1$ pins down marginal costs c_H and c_L, $w = 1$ also pins down the prices of final-good varieties. This determines prices in all three markets.

B Proofs of Lemma 2, Theorem 2 and Corollary 1

We will appeal to figure 1 in these proofs so we change variables from φ to $\varphi^{1-\sigma}$. To this end, let $H(\varphi^{1-\sigma}, \lambda)$ be the cumulative distribution function induced by this change of variable applied to $G(\varphi, \lambda)$. Let $\lambda(\varphi^{1-\sigma})$ be the λ that satisfies $\pi(\varphi^{1-\sigma}, \lambda) = 0$ in the region of figure 1 where $\varphi^{1-\sigma} > (\Lambda_L B)/(f + f_M)$. Then the free-entry condition is

$$
\begin{aligned}
f_E = & \underbrace{\int_0^{\frac{\Lambda_L B}{f+f_M}} \int_1^{\frac{\Lambda_L}{\Lambda_H}} \pi(\varphi^{1-\sigma},\lambda)\, H(d\varphi^{1-\sigma},d\lambda)}_{\mathcal{R}_0} + \underbrace{\int_0^{\frac{\Lambda_L B}{f+f_M}} \int_{\frac{\Lambda_L}{\Lambda_H}}^{\infty} \pi(\varphi^{1-\sigma},\lambda)\, H(d\varphi^{1-\sigma},d\lambda)}_{\mathcal{R}_a \text{ left}} \\
& + \underbrace{\int_{\frac{\Lambda_L B}{f+f_M}}^{\infty} \int_{\lambda(\varphi^{1-\sigma})}^{\infty} \pi(\varphi^{1-\sigma},\lambda)\, H(d\varphi^{1-\sigma},d\lambda)}_{\mathcal{R}_a \text{ right}} .
\end{aligned} \tag{D.1}
$$

We next differentiate this equation with respect to λ_L in two stages: (1) differentiation via the limits of integration and (2) differentiation of the integrand. In what follows we hold a^* constant because (1) if $a^* > 0$ then from the envelope theorem we can ignore changes in a^* and (2) if $a^* = 0$ then a^* stays at 0.

1. The derivative of the RHS of equation (D.1) with respect to $(\Lambda_L B)/(f + f_M)$:

$$
\int_1^{\frac{\Lambda_L}{\Lambda_H}} \pi(\tfrac{\Lambda_L B}{f+f_M},\lambda)\, H(\tfrac{\Lambda_L B}{f+f_M},d\lambda) + \int_{\frac{\Lambda_L}{\Lambda_H}}^{\infty} \pi(\tfrac{\Lambda_L B}{f+f_M},\lambda)\, H(\tfrac{\Lambda_L B}{f+f_M},d\lambda) - \int_{\lambda(\frac{\Lambda_L B}{f+f_M})}^{\infty} \pi(\tfrac{\Lambda_L B}{f+f_M},\lambda)\, H(\tfrac{\Lambda_L B}{f+f_M},d\lambda) .
$$

From the vertical line in figure 1, the integrand of the first term is $\pi_L = 0$ so the term vanishes. The second and third terms also vanish because they sum to zero. This follows from figure 1, which shows that $\lambda(\frac{\Lambda_L B}{f+f_M}) = \Lambda_L/\Lambda_H$.

2. The derivative of the RHS of (D.1) with respect to Λ_L/Λ_H:

$$
\int_0^{\frac{\Lambda_L B}{f+f_M}} \pi(\varphi^{1-\sigma},\tfrac{\Lambda_L}{\Lambda_H})\, H(d\varphi^{1-\sigma},\tfrac{\Lambda_L}{\Lambda_H}) - \int_0^{\frac{\Lambda_L B}{f+f_M}} \pi(\varphi^{1-\sigma},\tfrac{\Lambda_L}{\Lambda_H})\, H(d\varphi^{1-\sigma},\tfrac{\Lambda_L}{\Lambda_H}) .
$$

But this equals zero.

3. Equation (D.1) simplified: Combining the two results above, when we differentiate the free-entry condition with respect to λ_L we can ignore the limits of integration and bring the derivative inside the integrals of (D.1). Thus

$$
\int_{(\varphi,\lambda)} \left\{ \left. \frac{\partial\pi(\varphi,\lambda)}{\partial\lambda_L}\right|_{dB=0} + \frac{\partial\pi(\varphi,\lambda)}{\partial B}\frac{\partial B}{\partial\lambda_L} \right\} dH(\varphi,\lambda) = 0 \tag{D.2}
$$

or

$$
\frac{dB}{d\lambda_L} = -\left\{ \int_{(\varphi,\lambda)} \left. \frac{\partial\pi(\varphi,\lambda)}{\partial\lambda_L}\right|_{dB=0} dH(\varphi,\lambda) \right\} \Big/ \left\{ \int_{(\varphi,\lambda)} \frac{\partial\pi(\varphi,\lambda)}{\partial B}\, dH(\varphi,\lambda) \right\} \tag{D.3}
$$

B.1 Proof of Lemma 2

Depending on the value of (φ,λ), $\pi(\varphi,\lambda)$ equals 0, π_L, or $-a^* + p(a^*)\pi_H + \big(1 - p(a^*)\big)\pi_L$. The numerator of (D.3) is positive because an increase in λ_L increases Λ_k (Shephard's lemma), which increases π_k, $k = L, H$. Therefore π rises and the integrand is positive. The denominator of (D.3) is positive because an increase in B directly increases both π_L and π_H and so also increases π. It follows that $\partial B/\partial\lambda_L < 0$, as required.

B.2 Proof of Theorem 2

π is linear in $\Lambda_L B$ and $\Lambda_H B$. So we can rewrite the integrand in equation (D.2) in the form

$$d_1(\varphi,\lambda)\frac{\partial \Lambda_L B}{\partial \lambda_L} + d_2(\varphi,\lambda)\frac{\partial \Lambda_H B}{\partial \lambda_L}$$

where $d_1(\varphi,\lambda)$ and $d_2(\varphi,\lambda)$ are strictly positive on $\mathcal{R}_0 \cup \mathcal{R}_a$ and so are strictly positive with positive probability. Let $D_1 > 0$ and $D_2 > 0$ be their integrals. Then we can rewrite equation (D.2) as $D_1\partial(\Lambda_L B)/\partial\lambda_L + D_2\partial(\Lambda_H B)/\partial\lambda_L = 0$ or $\partial(\Lambda_H B)/\partial\lambda_L = -(D_1/D_2)\partial(\Lambda_L B)/\partial\lambda_L$ or

$$\frac{\partial \ln \Lambda_H B}{\partial \ln \lambda_L} = -\left(\frac{D_1}{D_2}\frac{\Lambda_L}{\Lambda_H}\right)\frac{\partial \ln \Lambda_L B}{\partial \ln \lambda_L}\,. \tag{D.4}$$

Since Shephard's lemma (lemma 1) and assumption 2 imply $\partial \ln \Lambda_L/\partial \ln \lambda_L > \partial \ln \Lambda_H/\partial \ln \lambda_L$, (D.4) implies

$$\frac{\partial \ln \Lambda_L B}{\partial \ln \lambda_L} > 0 > \frac{\partial \ln \Lambda_H B}{\partial \ln \lambda_L} \qquad \text{which implies} \qquad \frac{\partial \Lambda_L B}{\partial \lambda_L} > 0 > \frac{\partial \Lambda_H B}{\partial \lambda_L}\,.$$

By equations (5)–(6), this implies

$$\frac{\partial(\pi_H(\varphi,\lambda) - \pi_L(\varphi,\lambda))}{\partial \lambda_L} < 0\,.$$

By equation (8) with $(\varphi,\lambda) \in \mathcal{R}_a$ (meaning $a^* > 0$), this implies $\partial a^*(\varphi,\lambda)/\partial\lambda_L < 0$. This establishes equation (11) and hence theorem 2.

B.3 Proof of Corollary 1

Consider the equation in corollary 1. The left side is the total impact and is negative by theorem 2. On the right side, the direct effect may be negative or positive by theorem 1. Consider the last term of the equation. Since the firm has strictly positive innovation, figure 1 shows that $\lambda > \Lambda_L/\Lambda_H$ or $\lambda\Lambda_H - \Lambda_L > 0$. From the first-order condition (equation 8), $a^* = h(\pi_H - \pi_L)$ with $h' > 0$. So the sign of $\partial a^*/\partial B$ is the same as the sign of $\partial(\pi_H - \pi_L)/\partial B$. But since the firm innovates, $\pi_H - \pi_L = (\lambda\Lambda_H - \Lambda_L)\varphi^{1-\sigma}B > 0$, which is strictly increasing in B. Hence $\partial a^*/\partial B > 0$. From lemma 2, $\partial B/\partial\lambda_L < 0$. Hence the last term of the equation in corollary 1, $(\partial a^*/\partial B)\cdot(\partial B/\partial\lambda_L)$, is negative.

C Data Details

Countries: Poor countries are CHN (including HKG and MAC), MEX, as well as KHM, COL, FJI, GTM, HND, IND, IDN, JAM, KEN, MKD, MYS, MRT, PAN, PER, PHL, SRB, ZAF, LKA, SDN, THA, TUR, and VNM. Rich countries are USA as well as AUS, AUT, BEL, DNK, FIN, FRA, DEU, GRC, ISL, IRL, ISR, ITA, JPN, KOR, NLD, NZL, PRT, ESP, SWE, and GBR.

Magnitudes and Δm_{Lf}: Our key independent variable is the 2002–2011 change in low-wage offshoring as a fraction of total material costs. Since material costs include many inputs that are necessarily sourced domestically such as energy, it can be hard to get a handle on what level we should expect for offshoring as a fraction of material costs. A good benchmark is the fraction of Canadian material costs that were sourced from high-wage countries (read US). For firms in our sample this averaged 0.201 in 2002 and 0.181 in 2011. Given how integrated Canadian manufacturing is with US manufacturing (see section 5) this is an upper bound for what to expect. The low-wage country counterpart, dominated by Mexico and especially China, averaged just 0.018 in 2002 (a tenth of the US) but rose dramatically to 0.046 in 2011 (a quarter of the US). The change was 0.027 = 0.046 – 0.018. Most of this rise reflects a substitution away from the US and towards China and Mexico. While we do not have the firm-level data for subsequent years, by 2022 we estimate that low-wage offshoring had risen to 0.068. Thus the 2002–2022 change is 0.50 = 0.068 – 0.018.[39]

R&D Lumpiness: A small number of R&D zeros are due to the lumpiness of R&D expenditures and we deal with this by working with the averages $rd_{f,2002} \equiv (rd'_{f,2001}+rd'_{f,2002})/2$ and $rd_{f,2011} \equiv (rd'_{f,2010}+rd'_{f,2011})/2$ where rd'_{ft} is R&D expenditures in year t. Thus, when we refer to $rd_{f,2002}$ and $rd_{f,2011}$ in the text we mean the 2001–2002 and 2010–2011 averages, respectively. Using averages makes little difference.

D Variability of Shocks

We start by showing that we have a large number of shocks and these display substantial variability. Our data have 6,024 firms who are subject to 3,821 HS6-level shocks. The mean (interquartile range) across 3,821 shocks is 0.100 (0.283) for $\Delta\lambda_g$ and 0.156 (0.186) for $p_{Lg,2002}$. The average firm imports 34 HS6 intermediate inputs. Recall that s_{fg} is the share of HS6 good g in firm f's imports. It is a measure of the importance of shock g for firm f. The largest s_{fg} is a small 0.016 (only 1.6% of a firm's imports) or half the size of in the Autor et al. (2013) dataset. This is one indicator of substantial variability. Borusyak et al. examine the variability of shocks by examining the average importance of each shock, $s_g = \sum_f s_{fg}$, and calculating its Herfindahl Index $\sum_g s_g^2$. Its inverse measures the 'effective' number of observations. In our data, the inverse equals 546, which is three times larger than in the Autor et al. (2013) dataset. Large s_g are a sign that shocks are heavily concentrated on just a few HS6 goods so we would like the largest s_g to be small. In our data it is 0.0156, half the size of what appears in the Autor et al. (2013) dataset.[40]

E Entry and Exit

Our use of a balanced panel precludes empirical analysis of entry and exit. In our model, low-wage offshoring leads to the exit of some R&D performers, further reducing R&D. Low-wage offshoring also leads to the entry of some non-R&D performers, again further reducing R&D. Thus, entry and exit are additional channels through which low-wage offshoring negatively affects R&D. By ignoring them, we have understated the negative impacts of offshoring.

[39] The UN WITS database reports Canadian imports from China and Mexico by end-use category 'intermediates.' Between 2011 and 2022 this rose by a factor of 1.857. Statistics Canada reports cost of materials and supplies for manufacturing. Between 2011 and 2022 this rose by a factor of 1.265. Hence imports of intermediates scaled by material costs from China and Mexico (M_L) rose on average by 1.857/1.265 = 1.468. Thus, $M_{L,2022} = 1.468 \cdot M_{L,2011} = 1.468 \cdot 0.046 = 0.068$.

[40] Following table 2 in Borusyak et al. we decomposed variation in the HS6 shocks into 2-, 3- and 4-digit NAICS components and a residual using a random-effects, hierarchical model. For $p_{Lg,2002}$, the respective inter-class correlations are 18%, 14%, 15% and 53%, respectively. Further, the 2- and 3-digit NAICS components are statistically insignificant. Likewise for $\Delta\lambda_g$ where the residual (within NAICS4) is much larger. This supports the assumption that shocks are mean-independent across NAICS3 clusters, so it will be sufficient to cluster standard errors at the level of NAICS3 groups.