

Working Paper Series / Cahiers de recherche

February 2019, février

Processing Immigration Shocks: Firm Responses on the Innovation Margin

Rowena Gray (University of California, Merced)
Giulia Montresor (University of Essex)
Greg C. Wright (University of California, Merced)



Productivity Partnership

Partenariat productivité

productivitypartnership.ca

partenariatproductivite.ca

PROCESSING IMMIGRATION SHOCKS: FIRM RESPONSES ON THE INNOVATION MARGIN*

Rowena Gray[†], Giulia Montresor[‡] and Greg C. Wright[§]

February 14, 2019

Abstract

The extent to which firms respond to labor supply shocks has important implications for local and national economies. We exploit firm-level panel data on product and process innovation activities in the United Kingdom and find that the large, unanticipated, low-skill labor supply (immigration) shock generated by the 2004 expansion of the European Union to Eastern European countries increased process innovation and reduced product innovation. This implies that the innovation response to labor supply shocks may be more nuanced than the previous literature, focused mainly on patents, has suggested. Both of these effects are increasing in the low-skill intensity of firm production. In addition, the reduction in product innovation is lessened for firms whose output is sold locally, which is consistent with a demand side effect generated by the labor supply shock. We present a model that illustrates the channels through which firms may respond to labor supply shocks and find that both reduced form and instrumental variables results are mostly consistent with the model's predictions.

Key Words: Product Innovation, Process Innovation, Immigration, Labor Supply Shock

EL Codes: J23, J61, F22, O31, O33

*We thank many seminar participants, Laura Giuliano and Ethan Lewis for helpful comments. Montresor acknowledges the support of an Economic and Social Research Council PhD Studentship. We thank the University of Essex for facilitating continual access to the secure data and providing Greg Wright with a visitor travel grant.

[†]Gray: Department of Economics, University of California Merced. Email: rowenaegray@gmail.com.

[‡]Montresor: Department of Economics, University of Essex. Email: gmontr@essex.ac.uk.

[§]Wright: Department of Economics, University of California, Merced. Email: gwright4@ucmerced.edu.

1 Introduction

The impact of immigration on local economies is the subject of a large literature. Among other important findings, there is strong evidence that even large inflows of immigrant workers produce little impact on local employment rates and wages (see Card [13] for a discussion). Several mechanisms have been suggested to explain this. For instance, it has been shown that native and immigrant workers in many cases have complementary skills, even within low-education skill categories, which leads to productivity gains when these workers are used together (Peri and Sparber [41]). At the firm level, there is evidence that firms alter their production methods to use the now more abundant factor more intensively (e.g., Dustmann and Glitz [21] and Lewis [36]) while, possibly at the same time, adjusting their capital stock or adopting new technologies in response to the labor supply shock (e.g., Lewis [35], Lewis [36], Lafortune et al. [34] and Ottaviano and Peri [39]), both of which may mitigate any local wage and employment effects. More generally, Acemoglu [1] argues that firms will respond to changing skill supplies and premia by re-optimizing over the technologies used. In this paper, we explore two channels of firm response to labor supply shocks, namely, firm organizational changes arising from process innovations and changes in firm product scope due to product innovations, both of which may impact the distribution of output and employment within and across firms, with subsequent welfare consequences for workers.

There is little extant literature that separately relates labor supply shocks to process or product innovations, though Lewis [36] outlines a range of potential firm responses, some of which are consistent with the mechanisms we present here.¹ This gap in the literature is somewhat surprising, as organizational changes (process innovations) have been shown to be a key aspect of the firm response to technology adoption² and international trade,³ suggesting that labor supply shocks may also induce these types of responses. Here we provide evidence to this effect.⁴ In addition, we provide new evidence on the role that labor supply shocks play in firms' decisions regarding optimal product mix and scope (product innovations).

Generally speaking, the welfare consequences of process and product innovations are likely to be of first order importance, via well-known channels. Falling production costs due to process innovations will typically lead to price reductions and corresponding welfare gains. Relatedly, Cortes [15] shows that immigrants reduce prices of immigrant-intensive output, though in that case it is by reducing wages, and hence production costs. The channel we develop in this paper

¹The closest paper to ours is Maré et al. [37] who study the relationship between innovation and immigration in New Zealand, finding that the relationship between immigration inflows and innovation outcomes is a function of firm characteristics. We also find that firm responses are mediated by firm characteristics, as we discuss further in our results section.

²See Markus and Robey [38] for an early discussion and Bloom et al. [10], Bloom et al. [11] and Gaggli and Wright [27] more recently.

³See, e.g., Antras et al. [3] and Antràs and Rossi-Hansberg [4].

⁴For instance, a large, low-skill-intensive firm, UPS Parcel Delivery, recently stated that they would increase automation in their UK operations due to the expected UK exit from the European Union, an instance of (potential) process innovation in direct response to a shock to labor supply. See the Financial Times article November 2, 2016: <https://www.ft.com/content/e514de74-a0e3-11e6-86d5-4e36b35c3550>.

is an additional potential mechanism through which production costs and prices may fall due to increased immigration, in this case via increased process innovation. At the same time, increased process innovation may straightforwardly raise worker incomes by increasing firm productivity (e.g., see Huergo and Jaumandreu [31] or Bloom et al. [9]). And finally, product innovations should lead to a fall in the price index, and will thus increase welfare, by expanding the range of available product varieties (e.g., see Feenstra [25] or Eizenberg [23]).

At the level of the firm, Table 1 presents what we refer to as “innovation premia”, which are the estimates from a simple Ordinary Least Squares (OLS) regression of firm outcomes on an indicator for either product or process innovation (Section 4 describes the dataset). We alternately report estimates with and without firm fixed effects. Overall, we see that both product and process innovation are associated with greater firm revenue and employment, and greater overall production efficiency (labor productivity). Intuitively, we see that product innovation is associated with greater output and employment relative to process innovation, but process innovation is associated with relatively greater firm efficiency. Though these estimates are clearly not identified, they are suggestive of an important role for innovation in firm outcomes, consistent with a large literature. We return to these correlations in our discussion of the empirical results in Section 6.

We begin by presenting and discussing several stylized facts with respect to UK innovation, which we then use to motivate a model in which heterogeneous firms produce an endogenous set of branded varieties and employ both low- and high-skill workers. The firms’ product and process innovation decisions are made in order to achieve their optimal product scope and their optimal production structure, respectively.⁵ In our comparative statics exercise we focus specifically on a low-skill labor supply shock, first showing that firms increase process innovation in response. Furthermore, we show that firms that employ low-skill workers relatively more intensively will engage in relatively more process innovation in response to the shock, due to the fact that they reap greater profits from reorganizing their production structures to take advantage of the now-more-abundant low-skill labor.

We then show that product innovation could either increase or decrease due to the labor supply shock, depending on the relative magnitudes of each of three channels. First, the increase in the local labor supply also increases the local demand for goods and services, which allows firms to sell more of each product but also increases competition in the product market. We show that this increased competition leads firms to *reduce* their product scope, a result that is similar to the product “cannibalization” effect highlighted in the trade literature.⁶ Second, we show that to the extent that new product innovations are generated by high-skill workers, a low-skill labor supply shock will *reduce* product innovation when high- and low-skill workers are imperfectly substitutable, as the increase in low-skill labor supply raises the relative high-skill wage. Finally, the increase in low-skill

⁵This is how we think of process innovation throughout the model and empirics, i.e., as a set of firm-level organization changes enacted to achieve a more optimal production structure. We provide basic empirical evidence on “what process innovation is” in Section 4.

⁶See Eckel and Neary [22], Dhingra [16] or Hottman et al. [30].

labor supply reduces firm production costs in the short run (due to a falling low-skill wage), which has an ambiguous effect on product scope. This is because the fall in production costs increases competition in the product market, which again leads to within-firm product cannibalization (i.e., a reduction in product scope), but also makes all products more profitable to produce, which promotes an increase in product scope. We conclude that, since the net effect depends on the relative values of model parameters, the effect of a labor supply shock on product innovation is ultimately an empirical question.

We then bring the model’s predictions to UK data by exploiting the expansion of the European Union (EU) to Eastern Europe in 2004 as a differential, and large, shock to the supply of low-skill labor across UK local labor markets. This large inflow of immigrants to the UK was mostly unanticipated since, historically, the UK was a low-immigration country and expert predictions of the potential inflows due to the expansion were quite small. Using firm-level panel data on product and process innovation activities, we estimate several difference-in-differences specifications that produce consistent results, where our preferred specifications instrument for the change in the share of immigrants in each labor market (i.e., the intensity of treatment) using a shift-share measure that exploits network-driven migration from baseline distributions of immigrants across the UK in either 1991 or 2001. We first find that the immigration shock increased process innovation, on average. Noting that firms are likely to respond to supply shocks very differently, we then explore heterogeneity in the response. In fact, 77 percent of the variation in total UK employment growth over the period 2004-2010 occurred *between* firms (see Appendix A), suggesting that firms likely respond to shocks in heterogeneous ways. Indeed, consistent with the model predictions we find that the response was increasing in firms’ low-skill production intensity as well as firm size. We also find that product innovation *fell* in response to the migration, an effect that is greater within low-skill intensive firms but smaller for firms whose output is sold locally. We interpret this last finding as evidence on the importance of a demand side effect, though in fact its direction goes against the prediction of the model. We discuss potential reasons for this result in Section 6.1.

Existing research on the impact of labor supply shocks on innovation has typically focused on the impact of high-skill immigrants on patenting and knowledge creation more generally. For instance, Stuen et al. [43] exploit plausibly exogenous variation in the supply of foreign doctoral students in science to measure their impact on knowledge creation in the US, finding a large, positive and statistically significant impact. Hunt [32] also looks at the impact of immigrant students and finds that they patent at twice the rate of natives and are more concentrated in research-intensive fields such as science and engineering, without causing large crowding out of native innovation activities. Kerr [33] provides a more comprehensive review of studies looking at skilled immigration and innovation outcomes as proxied by patenting and firm starts. For the US, immigrants are found to play an important role in maintaining the country’s position as the technological leader in many fields, and particularly across STEM fields, with Chinese and Indian innovators being especially important in these areas. Another strand of work argues that diversity among high-skill workers

leads to higher levels of productivity and innovation, because diversity implies the interaction of complementary workers (see, for example, Ozgen et al. [40]). Our model also assumes that product innovation requires high-skill labor as the main input but our focus throughout the paper is the impact of a low-skill labor supply shock.

The paper is organized as follows. Section 2 presents some stylized facts. In Section 3 we jointly model the firm’s process and product innovation choice in the face of a labor supply shock. Section 4 describes the data. Section 5 introduces the empirical specifications and identification strategy. Section 6 discusses the results and Section 7 concludes.

2 Stylized Facts

2.1 EU8 Immigration to the UK as an Exogenous Shock

This paper exploits a large, policy-induced shock to the relative supply of low-skill labor across UK travel to work areas in the form of the expansion of the EU in 2004. Travel to work areas (TTWAs) are standardized UK local labor markets, a geographic unit developed by the Office of National Statistics (ONS). In short, these labor markets are defined in order to cover both metropolitan areas as well as their commuter suburbs.⁷ The expansion brought in eight Central and Eastern European countries: Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia and Slovenia, with the majority (75%) of UK immigrants from these countries over the 2004 to 2006 initial wave coming from Poland. We refer to these countries collectively as the EU8 countries. Though EU8 citizens immediately enjoyed free movement across EU countries, their access to most labor markets was restricted during a seven-year phase-in period. The exceptions were Ireland, Sweden and the UK who granted immediate access, the result of which was a large inflow of migrants into these countries. The UK restricted their access to benefits, so that migrants can be expected to be fully engaged in the labor market during this period. Blanchflower and Lawton [8] bear this out, showing that EU8 immigrants from 2004 to 2008 were 13 percentage points more likely to be working compared to natives and 5 percentage points more likely than pre-2004 immigrants from EU8 countries. Figure 2 depicts the long-run trend in immigration to the UK, indicating that 2004 was a significant turning point. In Figure 1 we see that this discontinuity is largely driven by the EU-accession-driven inflow of EU8 immigrants beginning in 2004. Indeed, in 2000 about 80 percent of migrants from EU8 countries lived in Germany and Austria, while after 2004 over 50 percent lived in the UK and Ireland (see Elsner [24]).

Most important for the purposes of our research design is that the magnitude of the inflow to the UK was largely unanticipated. Negotiations for the terms on which the new countries would enter the EU and enjoy its benefits, including full labor mobility, concluded only in December 2002 and the most highly publicized report at the time estimated that the net annual inflow from the

⁷Formally, the ONS defines a TTWA as a collection of wards for which “of the resident economically active population, at least 75% actually work in the area, and also, that of everyone working in the area, at least 75% actually live in the area”.




Figure 1: EU8 Immigration, 2004-2013. Source: ONS

new countries to the UK would be 5,000-13,000.⁸ At the time that the document was published, it was not known with certainty whether or not Germany would impose labor controls on the new accession countries, and so the authors estimated an extra 20,000-210,000 immigrants for Germany but emphasized that if Germany maintained labor controls then some of this expected flow might divert to the UK. The low anticipated flows for the UK were likely believable for UK firms, given the historically low immigrant inflows to the UK, particularly from these countries, and the stated preference of individuals in the new accession countries to move to locations closer to home both culturally and linguistically (Germany and Austria were the top destinations of choice as listed in the Home Office Report). Most commentary suggests the decision to maintain open borders immediately upon the new accession was taken solely by the UK government, with limited consultation from labor market actors, unlike the debates that occurred in other large European countries that maintained barriers to labor migrants which involved trades union and other labor partners. Such countries did not fully open their borders for labor migration until 2011, after our sample period closes.

Also important is the fact that the average hourly wage of EU8 migrants over the period 2004-2008 was far below that of the native population.⁹ According to Dustmann et al. [19] the average hourly wage over the period 2004-2009 for men from EU8 countries was £6.81 while it was £11.91 for native-born men. Blanchflower and Lawton [8] show that the most common occupations for EU8 workers up to 2008 were process operative and warehouse operative. This suggests that the EU8 expansion significantly changed the labor force composition in areas that received significant numbers of these migrants, skewing it more towards low-skill labor. To the extent that there is a fall

⁸See Dustmann et al. [18].

⁹This was despite their higher average education level, suggesting that either their origin country education was of low quality or was simply not highly relevant for jobs in the UK (see Dustmann et al. [19]).




Figure 2: Long-Run Trend in UK Immigration, 1975-2012. Source: ONS

in the average low-skill wage, this will have generated a productivity gain for firms who employed these workers, and relatively more so for firms who used low-skill labor relatively intensively, as we discuss further in the model.

2.2 UK Innovation

In Figure 3 we plot the unconditional change in the share of EU8 immigrants in a TTWA cell over the period 2004-2008 against the mean change in process (Panel A) and product (Panel B) innovation across firms over the period (see Section 4 for a description of the immigration and innovation datasets). The plots are clearly only suggestive, but they indicate a positive correlation between the group of relatively low-wage EU8 migrants and the extent of process innovation and a negative correlation between these migrants and the extent of product innovation. This suggests that process innovation may be induced by the availability of a new set of skills or, perhaps, simply by the increased availability of low-cost labor, which may induce a reorganization of production processes. On the other hand, product innovation is seemingly reduced by this inflow of workers, which potentially represents a welfare loss to consumers.

These facts help us to motivate the key assumptions of the model, specifically the extent to which the costs of product and process innovation are reliant on existing skill, as we discuss in Section 3 below. We then apply the formal predictions of the model and our identification strategy exploiting the unanticipated policy shock to relative labor supply to explore these relationships in more detail.




Figure 3: Change in Innovation vs Change in EU8 Immigrants Across UK Travel to Work Areas, 2004-08.

3 Model

We present a model in which the demand side closely follows the setup from Dhingra [16], who models single-factor, multiproduct firms in an international trade context. Our supply side adds an additional factor in a Constant Elasticity of Substitution (CES) production framework and derives results for the impact on product and process innovation in the case in which one of the factors increases in supply (as during a positive immigration shock). The model highlights the channels through which supply side shocks interact with firm product and process innovation decisions.

3.1 Consumers

There are M consumers in each local labor market who maximize utility over consumption of a homogeneous good and a differentiated good. Agent m consumes some amount of the homogeneous good along with some amount of each variety $i \in \Omega_j$ associated with brand $j \in \mathcal{J}$ of the differentiated good. Specifically, preferences of agent m are given by:

$$U^m \equiv q_0^m + \alpha Q^m - \frac{\delta}{2} \int_j \int_i (q_{ij}^m)^2 di dj - \frac{\eta}{2} \int_j (q_j^m)^2 dj - \frac{\psi}{2} (Q^m)^2$$

where q_0 represents consumption of the homogeneous good, $q_j^m \equiv \int_i q_{ij}^m di$ is the agent's consumption of brand j varieties, $Q^m \equiv \int_j q_j^m dj$ is total consumption of all varieties across all brands, and α, δ, η and ψ are constants. Consumers maximize this utility subject to their budget constraint, given by $q_0^m + \int_j p q_j^m dj = I^m$, where I^m is agent m 's income and p is the price of variety i of brand j where $p_{00} = 1$ is the numeraire good. We further assume that $q_0^m > 0$ and that all agents are identical. Maximizing the utility function and aggregating the resulting individual demand functions across

all consumers, we get the following linear inverse demand for variety i of brand j :

$$p = \tilde{\alpha} - \frac{1}{M} \left(\delta q_{ij}^m + \eta q_j^m \right) \quad (1)$$

where $\tilde{\alpha} \equiv \alpha - \psi Q^m/M$ reflects demand conditions the firm takes as given.

The linear demand system (1) is useful, in part, because it is consistent with the empirical findings of Hottman et al. [30] who show that variation in product scope can explain a substantial portion of variation in sales across US firms. In addition, this demand system generates product cannibalization, a mechanism described by Dhingra [16] and one that Hottman et al. [30] find to be important in explaining firms' response to demand shocks. In short, and as we describe in detail below, cannibalization implies that each additional product produced by a firm both generates additional firm profits while also reducing the demand for the firm's existing products, with equilibrium determined by the balance of the two forces. This mechanism ultimately provides a tractable (and empirically relevant) condition to pin down the range of products produced by each firm, as we will show.

3.2 Firms

Each firm j is associated with a brand, and may supply multiple varieties within the brand to its local labor market. Throughout the analysis we focus on outcomes associated with a single market – i.e., we focus on the partial equilibrium – though there are in principle many markets and in the empirics we will exploit variation across multiple markets. There is free entry in the differentiated goods industry and, after paying a fixed entry cost, f , firms can enter and produce each variety i at marginal cost c . The firm's production function combines two labor types, high-skill and low-skill labor. An important feature of the model is that the firm can choose from an array of production methods, conditional on its given underlying production structure, and these differ in their relative efficiency of use of the inputs. When the firm adjusts the relative efficiency of its inputs we consider this to be *process innovation*.

The idea is that firms may respond to a shock to the relative labor supply not only by using labor types in different proportions, but also by altering their production methods to use the now-more-abundant factor more efficiently. Formally, the firm takes local factor prices as given and chooses from a continuous menu of production technologies. Beyond this, we assume a fixed heterogeneity in the intensity of use of labor inputs across firms. As a result, while the firm is able to adjust the relative efficiency of its inputs, it is simultaneously constrained by the unique, and fixed, production structure required to make its particular products.

Finally, apart from endogenously choosing the efficiency of its factors, the firm also endogenously chooses its optimal product (variety) scope, which we refer to as *product innovation*. As we will show, product innovation will, in part, depend on the firm's choice of process innovation, and each type of innovation will independently respond to labor supply shocks in the firm's local market.

Production. Having paid the fixed entry cost, the firm's variety-specific production technology is given by the following production function:

$$Y_{ij} = [\beta_{ij}(A_{ijL}L_{ij})^\rho + (1 - \beta_{ij})(A_{ijS}S_{ij})^\rho]^{1/\rho} \quad (2)$$

where L and S are low-skill and high-skill labor inputs, the efficiency parameters A augment each factor (and will become choice variables later on), and the elasticity parameter $\rho \equiv \frac{\sigma-1}{\sigma} > 0$. The terms β_{ij} and $1 - \beta_{ij}$ are exogenous, variety-specific technology terms that define the fixed input proportions firms are constrained to use to produce their varieties. This feature reflects the fact that the factor content of output is to some degree determined by the nature of the product being produced, and is therefore to some extent outside of the firm's control (at least in the short run).

In order to more flexibly define the notion of process innovation later on, we do not explicitly incorporate capital in the production function. There are two primary reasons: first, many examples of process innovation combine organizational changes with investments in capital, and it is more tractable to consider these jointly as an increase in one of the efficiency variables, A . Second, process innovation may be, at times, skill-biased and, at other times, unskill-biased. An example of the former is the incorporation of computer-assisted design software for product development (which may augment the productivity of engineers), while an example of the latter is the adoption of GPS systems for product delivery (which may augment the productivity of truck drivers). The production function, (2), again allows us to flexibly model these as different types of investments in factor efficiency.¹⁰

The production function, (2), indicates that the firm is constrained in its production process – reflected in the fixed β s – and at the same time has a degree of flexibility in that it can choose both the relative quantities of factors employed as well as the relative efficiency of its inputs, A_{ijL} , A_{ijS} . We also note that when varieties are symmetric in production there is no need for firm-product subscripts, and so we drop these subscripts henceforth.

Given the production function, (2), the cost minimizing choice of inputs is given by the usual first-order conditions (FOC) which equate the (exogenously determined, from the firm's point of view) wage paid to each factor with its marginal productivity. Formally, relative factor demand within a firm is given by:

$$\frac{L}{S} = \left[\frac{w_L(1-\beta)}{w_S\beta} \left(\frac{A_S}{A_L} \right)^\rho \right]^{1/(\rho-1)} \quad (3)$$

where, for the reasons noted above, we have dropped the firm-product subscripts.

When relative wages change, perhaps due to an increase in the local supply of one factor, the firm responds by increasing its relative use of that factor, in order to reduce the marginal productivity

¹⁰An alternative would be to combine each labor type with a capital type in a CES combination, with each combination then combined in an upper CES nest. This would give qualitatively similar results in a more complex setting.

of the factor and bring it back in line with its wage (conditional on the endogenous response of the efficiency terms).

Unit Costs. It is useful from this point on to work with the firm's unit cost function, which incorporates the firm's optimally chosen factor quantities, reflected in (3). Formally, minimizing factor costs subject to (2) we obtain the unit cost, c , associated with production of any firm (brand) variety, which is given by:

$$c = \left[(\beta)^\sigma \left(\frac{w_L}{A_L} \right)^{1-\sigma} + (1-\beta)^\sigma \left(\frac{w_S}{A_S} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \quad (4)$$

where w_l are factor prices that the firm takes as given, with $l \in (L, S)$, and the terms A_l and β are the endogenous and exogenous technology terms, respectively.

Process Innovation. We define process innovation to be a shift toward a new, more efficient production function by the firm. Specifically, we assume that any outward adjustment of the technology frontier requires expenditure by the firm. Formally, we assume that $A_l \equiv \tilde{A}_l(1 + \kappa_l)$, where the firm can invest in process $\kappa_l \in [0, \infty)$ in order to increase the efficiency of factor l , where *higher* levels of κ are associated with *lower* unit costs. \tilde{A}_l is then the firm's baseline factor efficiency. The firm can increase the efficiency of one of its factors by investing in process innovation at a rate r_l , so that expenditure on process innovation is given by $r_l \kappa_l$.¹¹ Here we note that the cost of process innovation is factor-neutral, such that a change in the supply of either factor will induce process innovation. This assumption is made in light of the stylized facts in section 2.2 that indicate a correlation between variation in the low-skill supply in an area and process innovation. This suggests that high-skill labor is not a pre-requisite condition for this type of innovation, though this is only a suggestive finding and is one we will explore more carefully in the empirics.

Product Innovation. We assume that the firm chooses its optimal product scope, h , producing an additional variety at a cost $r_h w_S$. The assumption is that product innovation – adding a new variety – requires payment of a variety-specific R&D cost at rate r_h , which is denominated in high-skill labor. For instance, adding a new product may require R&D expenditure on the wages of scientists and engineers, in contrast to process innovation which can perhaps be done by incurring

¹¹In a previous version we assumed that the firm faced a tradeoff in the extent to which it could engage in low-skill-biased process innovation versus high-skill-biased process innovation. In that case, we followed Caselli et al. [14] in modeling the shift as the choice of a new (A_L, A_S) pair in the available technology space. More formally, the firm's technology frontier – i.e., the choice set of available technologies – was given by:

$$(A_L)^\alpha + \eta(A_S)^\alpha \leq B$$

where η and α govern the tradeoff between the relative efficiency of each factor and B defines the height of the technology frontier, and is firm-specific. However, this produces nearly identical qualitative results, but with the size of the firm response to a shock governed also by the additional parameters associated with the above technological constraint. In this version we instead pursue the simpler case in which the firm faces no tradeoff with respect to performing either type of process innovation.

costs that are not dependent on the skill composition of the firm's workforce. As discussed above, this assumption is made in light of the stylized facts in section 2.2, which indicate that a relative reduction in high-skill labor is associated with a reduction in product innovation.

Profit Maximization. Given these costs, total firm profits can be written as:

$$\Pi = \int_0^h [p - c(A_L(\kappa_L), A_S(\kappa_S))] q \, di - \int_0^h (r_L \kappa_L - r_S \kappa_S - r_h w_S) di \quad (5)$$

where c is given by (4) and the integrals are taken across (symmetric) products within the firm. For tractability, we assume throughout that firms and varieties are identical except for firm-specific heterogeneity in the production technology – i.e., we assume that only β varies across firms and that varieties are identical within a firm. As a result, we can re-write (5) as:

$$\Pi = h \left\{ [p - c(\kappa_L, \kappa_S)] q - r_L \kappa_L - r_S \kappa_S - r_h w_S \right\} \equiv h\pi \quad (6)$$

where π is the profit associated with each variety produced by the firm and we now simply write marginal costs as a function of the κ 's. Note that since firms' costs differ – due to the heterogeneity in β – their prices, quantities, the level of investment in process innovation and the number of varieties produced by a firm will also differ.

3.3 Equilibrium and Comparative Statics

We first solve for optimal q . Maximizing firm profits, the FOC is $\frac{\partial \pi}{\partial q} = p - q \left(\frac{\delta}{M} + \frac{h\eta}{M} \right) - c(\kappa_L, \kappa_S) = 0$. Combining this with inverse demand, optimal firm output is therefore given by

$$q^* = \frac{M}{2(\delta + h\eta + 1)} (\tilde{\alpha} - c(\kappa_L, \kappa_S)) \quad (7)$$

The optimal values of low- and high-skill process innovation are then given by the profit-maximizing expenditure on each, i.e., $\{\kappa_L^*, \kappa_S^*\}$, while optimal product innovation is given by the profit maximizing product scope.

We are primarily interested in the comparative statics with respect to an increase in the low-skill labor supply in an area, and so that is what we focus on here. We focus on the associated cost function, (4), in which the endogenous choice of technique – i.e., the choice of κ_L – operates above and beyond the firm's adjustment of its relative use of factors.¹²

In the analysis that follows we will assume that, in the short run, w_L unambiguously falls when the supply of low-skill labor rises, and that relative factor adjustment within the firm, as described by (3), only partially mitigates the fall in the low-skill wage generated by the increased local supply of low-skill labor. In making this assumption, we are able to highlight firms' innovation responses as

¹²Of course, the cost function explicitly incorporates the firm's optimal choice of factors.

a mechanism that may subsequently put *additional upward pressure* on the relative low-skill wage, beyond that due to the firm's adjustment of its relative use of factors.

Optimal Product Innovation. The FOC with respect to the firm's choice of number of varieties is pinned down by the linear demand, (1). As shown by Dhingra [16], the linear demand system causes new varieties to cannibalize the demand for existing varieties. As a result, the additional profit that the firm obtains due to an increase in product scope is countered by a decline in overall profits as demand for existing products falls. The balance of these forces pins down the optimal number of varieties, where the profit from the marginal variety is equal to the decline in aggregate profits due to cannibalization. This optimal product scope is given by the solution to the FOC, $\frac{\partial \Pi}{\partial h} = 0$, which is:

$$h^* = \frac{\pi^* M}{(q^*)^2 \eta} \quad (8)$$

where π^* are optimal profits.

Given this equilibrium condition, an increase in the low-skill labor supply in an area generates three primary effects on the product margin which are summarized in Proposition 1 below. First, differentiating (8) with respect to M , the size of the local market, we find that $\frac{\partial h^*}{\partial M} < 0$.¹³ This is the somewhat counter-intuitive result that is analogous to the trade context described in Eckel and Neary [22] and Dhingra [16]. In short, firms respond to the overall rise in demand by increasing output per product (q^*) while reducing the number of products (h^*). This is because the increase in market size leads to entry of new firms, and thus greater product market competition, and this shifts the demand intercept for any individual variety inward, thereby reducing its profitability.¹⁴ Firms adjust to this fall in profitability by reducing their product scope in order to relax within-firm, across-product competition. In other words, by lessening competition across their own product lines they raise overall profits, thus offsetting the reduction in profits due to the now-greater competition in the larger market.

Second, differentiating (8) with respect to the low-skill wage (which falls due to the increased supply of low-skill workers), we find that the sign of $\frac{\partial h^*}{\partial w_L}$ is ambiguous. This is because there are competing effects: on the one hand, production costs fall due to the low-skill labor supply shock, which makes production of all varieties more profitable, which then increases the equilibrium range of profitable varieties. On the other hand, the fall in unit production costs increases competition between firms, much like the effect due to increased market size described above. Because of this, the firm prefers to reduce its product scope and produce more output per product instead. The sign of $\frac{\partial h^*}{\partial w_L}$ ultimately depends on the relative magnitude of model parameters, as well as the size of the market (M).

¹³Note that q^* and π^* are functions of M .

¹⁴In Eckel and Neary [22] and Dhingra [16] the expansion in market size is due to international trade, rather than migration.

Finally, since low-skill labor and high-skill labor are imperfect substitutes, the increased supply of low-skill labor leads to an increase in the high-skill wage (high- and low-skill workers are “Q-complements”). Since the cost of product innovation is denominated in terms of the price of high-skill workers, this reduces the profitability of all products, and therefore reduces the optimal product scope. We summarize these findings in the following Proposition:

Proposition 1 (Product Innovation Response). *From (8), there are three channels through which a low-skill labor supply shock impacts optimal firm product scope:*

1. $\frac{\partial h^*}{\partial M} < 0$. *By increasing the size of the local market, and thus increasing product-market competition, a low-skill labor supply shock decreases the return to product innovation, thereby **reducing product scope**.*
2. $\frac{\partial h^*}{\partial w_L} > < 0$. *By reducing production costs, a low-skill labor supply shock increases the range of profitable varieties, but also incentivizes the firm to reduce product scope due to a competition effect. The sign of $\frac{\partial h^*}{\partial w_L}$ is therefore **ambiguous**.*
3. $\frac{\partial w_S}{\partial w_L} < 0$. *Due to the imperfect substitutability of high- and low-skill labor, a low-skill labor supply shock increases the cost of product innovation, thereby **reducing product scope**.*

The direction of the overall effect is therefore ambiguous, and hinges on the sign and magnitude of the second channel.

Optimal Process Innovation. Since the FOC for high- and low-skill process innovation are symmetric, we solve for the FOC for low-skill process innovation. We again note that product and process innovations are simultaneously determined and endogenous to one another. With this in mind, we calculate the FOC for process innovation and plug in the values for optimal q^* and optimal h^* determined above, which leads to the following implicit equilibrium condition, denoted $F(\kappa_L^*)$, for optimal process innovation, κ_L^* :

$$F(\kappa_L^*) \equiv \frac{M\beta^\sigma w_L^{\sigma-1}}{(1 + \kappa_L^*)\left(\delta - 1 + \frac{\pi(\kappa_L^*)M}{q(\kappa_L^*)^2}\right)} \left((p + \tilde{\alpha})c(\kappa_L^*)^\sigma - 2c(\kappa_L^*)^{\sigma+1} \right) - r_L = 0 \quad (9)$$

Implicit differentiation of equilibrium condition (9) with respect to the low-skill wage w_L then leads to the following result:

Proposition 2 (Process Innovation Response). *Following from (9), $\frac{\partial \kappa_L^*}{\partial w_L} < 0$ iff*

$$\frac{[\sigma c(\kappa_L^*)^{\sigma-1} \chi_1 - \sigma + 2] [p - c(\kappa_L^*) + q(\kappa_L^*)] + 2c(\kappa_L^*)^\sigma \chi_2}{w(1 + \kappa_L) [\sigma c(\kappa_L^*)^{(\sigma-1)(1-\sigma)} \chi_3 - \sigma + 1] [p - c(\kappa_L^*) + q(\kappa_L^*)] - 2c(\kappa_L^*)^\sigma \chi_4} > 0 \quad (10)$$

where κ_L^* is the firm's optimal investment in low-skill-biased process innovation and $\chi_1 \equiv \frac{\beta^\sigma w^{1-\sigma} \tilde{A}_L^{1+\sigma}}{1+\kappa_L^*}$, $\chi_2 \equiv \frac{\beta^\sigma w^{1-\sigma} \tilde{A}_L^{1+\sigma}}{(1+\kappa_L^*)^{\sigma-1}}$, $\chi_3 \equiv \frac{\beta^\sigma w^{1-\sigma} \tilde{A}_L^{\sigma-1}}{(1+\kappa_L^*)^{\sigma-1}}$ and $\chi_4 \equiv \frac{\beta^\sigma w^\sigma \tilde{A}_L^{\sigma-1}}{(1+\kappa_L^*)^{-\sigma}}$.

See Appendix B for proof. Thus, when the proposition holds, a rise in the low-skill labor supply induces firms to increase the efficiency of their low-skill workers via process innovation. The intuition for the proposition arises from a straightforward tension within the firm in the face of falling input costs (e.g., the low-skill wage in our case). On the one hand, the firm would like to engage in more process innovation which will raise output and profits. On the other hand, process innovation is costly and so reduces profits. Proposition 2 implicitly defines the optimal innovation response to a falling low-skill wage in light of this tension.¹⁵

We also note that the FOC with respect to κ_L , $\frac{\partial \pi}{\partial \kappa_L} = -q \frac{\partial c}{\partial \kappa_L} - r_L = 0$, indicates that optimal process innovation is increasing in firm output.¹⁶ We formalize this and a further implication in the following lemma:¹⁷

Lemma 1 (Role of Firm Size in Process Innovation). *Optimal process innovation is increasing in firm output. In addition, the process innovation response to a local labor supply shock is also increasing in firm output – i.e., $\frac{\partial \kappa_L^*}{\partial w_L \partial q} < 0$.*

Furthermore, since firms are heterogeneous in their production structures, their responses to the low-skill labor supply shock are also heterogeneous. Specifically, $\frac{\partial \kappa_L^*}{\partial w_L \partial \beta} < 0$, such that firms whose production is relatively intensive in low-skill labor increase their investments in process innovation relatively more. We summarize this result in the following lemma:

Lemma 2 (Role of Factor Intensity in Process Innovation). *The process innovation response to a local labor supply shock is increasing in the firm's intensity of use of the now more abundant factor – i.e., $\frac{\partial \kappa_L^*}{\partial w_L \partial \beta} < 0$.*

Proofs of Lemmas 1 and 2 can be found in Appendix C.

Lemma 2 is particularly important due to the fact that we cannot directly observe the factor bias of process innovation within a firm in the data. This is because firms simply report whether or not they have undergone process innovation, but do not report much detail about what was actually done. As a result, in linking the model to the empirics we can exploit Lemma 2, which tells us that the process innovation response by the firm will be designed to augment whatever factor

¹⁵The complexity of the condition arises due to the endogenous responses of product innovation and output.

¹⁶Note that we do not explore the effect of size heterogeneity on product innovation in the previous section, though we do explore its consequences for process innovation here. This is because, for product innovation, firm size affects each of the channels summarized in Proposition 1, and not in the same direction. As a result, the formal condition that describes the net effect has no easy intuition associated with it and allows no easy test.

¹⁷Note that the following lemma is reliant on the fact that the cost of process innovation is independent of firm size – i.e., a unit of investment reduces the *unit* costs associated with some factor. If process innovation did depend on firm size then the following results would hinge on the functional form of the relationship between cost and size.

is now more abundant. As such, although we cannot see what type of process innovation is being implemented, we can infer it by observing which firms engage in process innovation in response to the shock, since the magnitude of the process innovation response is a function of the initial relative factor intensity of the firm. To the extent that the data are consistent with these facts, this can serve as confirmation of the underlying logic of the model.

3.4 A Note on General Equilibrium

The above analysis focuses on a single local labor market but, of course, general equilibrium outcomes may differ somewhat from those derived in partial equilibrium. The primary differences between partial and general equilibrium¹⁸ are that both labor markets and output markets must clear at the national level in general equilibrium (assuming no international trade). This will have consequences for the magnitude of the effects due to a local labor supply shock as local areas interact with one another to clear national markets.

With respect to labor market clearing, we would expect that a relatively large labor supply shock in some area may induce *out-migration* of existing workers from an area, as the return to labor falls locally. This may affect the magnitude of the innovation responses derived above, though each of the results will still hold qualitatively. With respect to output markets, a so-called Rybczynski effect will reallocate production across local labor markets, an effect that will work in the opposite direction by encouraging *in-migration* to areas that see the largest, positive labor supply shocks (see Hanson and Slaughter [28]).

Thus, there may be spillovers across areas that we abstract from in the partial equilibrium described above, though the net effect of these spillovers on outcomes is ambiguous. We discuss the implications of these general equilibrium effects further in our empirics below.

4 Data

In Section 5 below we explore the innovation response of UK firms to a labor supply shock in their local labor market, which we define as a UK Travel to Work Area (TTWA).¹⁹ The variation in labor supply across the 243 UK TTWAs that we exploit in generating our stylized facts comes from the UK Quarterly Labour Force Survey (QLFS). The QLFS is a quarterly sample of workers that includes a variety of work-related and demographic information, including the worker's country of birth. Our identification strategies also exploit cross-sectional variation in the EU8 immigrant share of the population and economic outcomes from the 1991 Census and immigrant shares from the 1981 Census. Summary statistics for our primary variables of interest are reported in Table 2.

Our dependent variables exploit firm-level panel data on innovation activities from three waves of the Community Innovation Survey (CIS), covering the period 2002 to 2008.²⁰ The CIS is the

¹⁸Of course, depending on the complexity of the model there may be many general equilibrium consequences that we abstract from here.

¹⁹We use the 2001 ONS definition of a TTWA.

²⁰Note that we use the population-representative panel dimension of the data in part because our specifications use

primary source of information on innovation for the UK, and asks firms a range of questions about their research and development activities as well as the extent to which they have undertaken various types of organizational change during the previous three years. It is conducted every four years, such that we exploit survey responses regarding firms' innovation activities at three snapshots: 1) between 2002 and 2004 – the period prior to the EU8 accession – as well as 2) between 2004 and 2006 and 3) 2006 and 2008. The nature of the timing of the survey requires two comments. First, there is an overlapping year in each survey, however this is inconsequential given the binary nature of our outcome variables.²¹ For instance, if a firm reports product innovation for the 2002-2004 period, and then no product innovation for 2004-2006, we know that the firm engaged in product innovation in 2004 (and, of course, 2002-2003). This effectively means that the later two datasets reflect firm behavior in the final two years that they cover.

Second, we must address the fact that 2004 is assigned to the pre-period, even though one of our identification strategies exploits variation due to the EU enlargement that occurred on May 1st 2004. As a result, any response by firms from May through December of 2004 due to the immediate inflow of immigrants from EU8 countries will be allocated to our pre-period control group. We note that this will work against finding an effect due to the EU8 accession – i.e., it will bias our results towards zero. Figure 1 documents the trend in EU8 inflows beginning in 2004. We can see that there was indeed an immediate uptick in EU8 migration to the UK beginning in June, 2004, however the vast majority of the inflow occurred after December 2004. Furthermore, given that the government vastly underestimated the projected magnitude of the immigrant inflows, suggesting there was little anticipation of the magnitude, it is likely that any firm response to the labor supply shock lagged the EU expansion.

Finally, the CIS consists of a stratified sample of approximately 28,000 firms with more than 10 employees. For the period we are interested in, 2002-2008, the CIS has a panel dimension consisting of a subsample of approximately 8,500 firms, and this is the sample that we exploit in our baseline analysis (with fewer observations in some specifications).

The CIS asked the following questions, which allow us to construct our outcome measures: 1) During the last three years did your enterprise introduce new or significantly improved goods or services and 2) During the last three years did your enterprise introduce new or significantly improved methods of manufacturing or producing goods or services.²² It further asked for firms' spending on R&D, and the objectives of these innovation activities. These questions regarding whether firms actually did innovation may be a more direct measure than the traditional patent data used in the literature, which measure invention rather than innovation. Similarly, we would not want to rely on R&D expenditure entirely because not all expenditures will successfully lead to implementation of new products or processes.

firm fixed effects. The unbalanced dataset would drop firms observed in only a single year and the resulting sample would be unrepresentative.

²¹We also exploit continuous variables from the CIS in our interaction regressions, but in these cases we only use data from the pre-period survey – i.e., we do not rely on variation over time in the response.

²²This is paraphrased from the 2008 CIS documentation.

What is Process Innovation?

The notion of process innovation is typically taken to be one type of organizational change; specifically, it usually reflects the implementation of more sophisticated or appropriate production processes in order to increase efficiency. A canonical example, analyzed in Basker [6], is the introduction of barcode scanners at grocery stores in the 1970s and 1980s, which revolutionized many aspects of the retail sector. Reassuringly, this is also what respondents to the CIS have in mind. In Table 3 we present the coefficients and standard errors from an OLS regression across all firm observations in all years, in which the process innovation indicator is regressed on the response by firms to questions regarding the extent to which they made one of several organizational changes, as well as the extent to which they made investments in capital (column 1). The latter variable is included in order to determine whether process innovation is simply a proxy for capital investments which, as noted in the Introduction, have been explored in the context of immigration in other papers. Since capital investments may simply be a proxy for firm size, we also re-run the multivariate regression controlling for capital investments *per worker* (column 2), in order to capture the determinants of process innovation conditional on capital intensity.

As we can see from the table, the strongest (conditional) correlates with process innovation are “Improvements in Production Flexibility”, “Improvements in Production Capacity” and “Reduce Per Unit Costs”. Additionally, there is virtually no correlation with contemporaneous capital investment or capital intensity. Across both specifications “Reduce Per Unit Costs” is highly significant, which is consistent with a standard theoretical treatment of process innovation as a simple scaling of unit costs (see, e.g., Dhingra [16] or Duranton and Puga [17]).

Variable	Coefficient	Coefficient
Improve Product Quality	0.018*	0.016*
Improve Production Flexibility	0.055***	0.042*
Improve Production Capacity	0.073***	0.050**
Reduce Per Unit Costs	0.061***	0.079***
Improve Health and Safety	0.004	0.005
Increase Value Added	0.011	0.009
Capital Acquisition (millions £)	0.059	
Capital Acq. Per Worker		0.071
Model R ²	0.79	0.84

Note: Dependent variable is a firm-year process innovation indicator. Column 1 controls for total capital acquisition, while Column 2 controls for capital acquisition per worker. * p<0.05, ** p<0.01, *** p<0.001

Table 3: Correlates with Process Innovation

5 Specifications & Identification

In this section we link the main predictions of the model to formal empirical specifications. We also outline our identification strategy, describe the underlying source of identifying variation, and perform a range of tests of the identifying assumptions.

5.1 Baseline

In our baseline specification we focus on the results summarized in Propositions 1 and 2. To reiterate, Proposition 1 states that the direction of the product innovation response to a labor supply shock is ambiguous, while Proposition 2 defines the conditions under which process innovation will rise or fall in response to such a shock. Formally, we exploit the discontinuous inflow of immigrants arising from the 2004 EU8 expansion, described in Section 2 above. To exploit the discontinuity, we begin by considering the following difference-in-differences specification:

$$INN_{iat} = c + \beta_1 [POST_t \times EU8SHR_{a,2004}] + \alpha_t + \gamma_i + \epsilon_{iat} \quad (11)$$

where INN is one of the binary innovation measures of interest, associated with firm i located in TTWA a in period t ; $EU8SHR$ is $100 \times$ the share of EU8 immigrants in TTWA a in year 2004, where we multiply by 100 so that a one-unit change in the share is equivalent to a one percentage point change; $POST$ is an indicator equal to 1 for post-2004 periods and 0 for the 2002-2004 period; α_t and γ_i are period and firm fixed effects, respectively; and ϵ_{iat} is the residual obtained from projecting the untreated potential outcome (i.e., when the treatment is zero) on the control variables.

The differential extent of the “treatment” is defined by the cross-sectional variation in the share of EU8 immigrants in a TTWA at the beginning of the period, 2004. This approach exploits a version of the “ethnic enclave” design commonly associated with Altonji and Card [2] and Card [12]. The idea is that immigrant groups tend to settle in locations in which their compatriots are already settled. As a result, the pre-existing distribution of a particular immigrant group – say, Hungarians – across locations will serve as a good predictor of the future pattern of Hungarian settlement in the U.K. The share of Hungarians in an area in 2004 should therefore serve as a useful predictor of Hungarian settlement between 2004 and 2008.

Figure 4 suggests that this is the case. The x-axis reports the share of immigrants in a TTWA from each of the EU8 origin countries in 2004, while the y-axis reports the percentage point rise in the share of immigrants from each country in a TTWA between 2004 and 2008. In other words, each observation represents a single TTWA-by-EU8-origin-country cell (243 TTWAs by 8 origin countries), and we relate the 2004 level to its subsequent growth. We see that areas with a higher share of immigrants from a particular EU8 country in 2004 on average saw a larger rise in immigration from that country relative to other immigrant groups in subsequent years. In effect, this represents a “first-stage” regression, and we note that the slope is statistically significant at the one

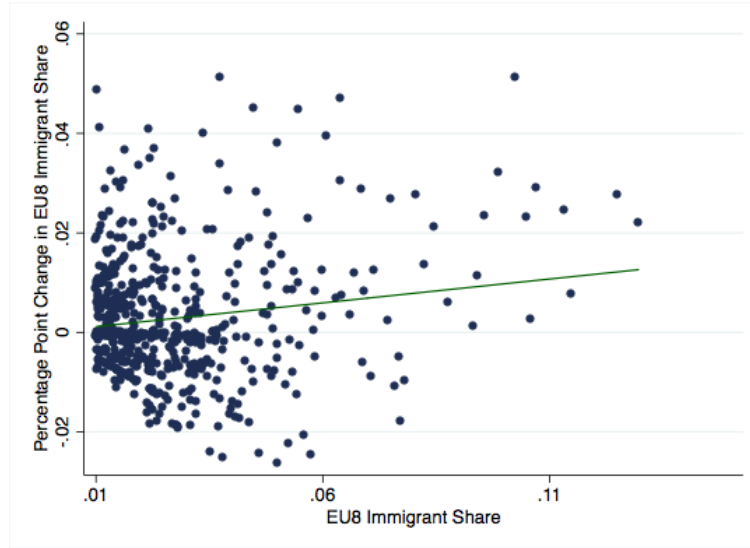


Figure 4: EU8 TTWA Immigrant Share, 2004 vs Percentage Point Change in Share, 2004-2008

percent level (p-value 0.003), with a large partial F-Stat of 41.

We estimate each version of (11) with both a linear probability model (LPM), our main specification, as well as a conditional logit (CL) specification, reported in an Appendix. Since in these specifications the treatment varies across TTWAs in the cross-section, we cluster standard errors at the TTWA level throughout. All regressions are weighted by TTWA employment though the qualitative findings are not sensitive to these weights, including the use of firm-size weights. In Section 6 below we discuss these estimates.

5.2 Identification Strategy

The exclusion restriction with respect to specification (11) is still likely to be violated, since innovation outcomes within a local labor market over the 2004-2008 period are likely to be driven by shocks that also drove immigrants to that area in the years leading up to 2004. Furthermore, a second challenge arises due to the fact that there may be spillovers (general equilibrium effects) across local labor markets in response to the treatment, as discussed in Section 3.4.

To deal with the first issue, we estimate specification (11) both using the 2004 distribution (as in (11)) and also using a lagged EU8 share variable, $EU8SHR_{a,1991}$ in place of the 2004 distribution, reflecting the share of EU8 immigrants in a TTWA in 1991.²³ The potential endogeneity problem now only arises if, for instance, a productivity shock that drove EU8 immigrants to an area in 1991 also influences firm-level innovation in that area over the period 2004 to 2008. In other words, if the hypothetical productivity shock is serially correlated (enough) then this may be the case, and there

²³Bell et al. [7] use the 2001 Census distribution of immigrants from EU8 countries to implement a similar Bartik-type procedure in an attempt to isolate the exogenous component of immigration across locations within the UK. Fogel and Peri [26] also apply a difference-in-differences strategy that is similar to ours. Specifically, their treatment is an interaction between an inflection point in which immigration rapidly increased and an indicator for whether a region was in the highest or lowest quartile of predicted inflows based on a historical refugee dispersal policy.

may be lingering endogeneity. We rely on the fact that 1991 was distant enough so that the shocks driving immigrants to particular TTWAs in 1991 are very likely to be uncorrelated with the shocks to innovation over the recent period, an assumption that we discuss and test in the next section.

We refer to these specifications as the reduced form regressions and add a third specification in which we instrument for $[POST_t \times EU8SHR_{a,2004}]$ with $[POST_t \times EU8SHR_{a,1991}]$. This specification has the benefit of producing coefficients whose magnitudes are readily interpreted as an immigrant enclave-driven Local Average Treatment Effect (LATE). We note that this IV approach assumes the exogeneity of the $POST_t$ indicator – i.e., that there are not coincident shocks unique to 2004 that also led to increased innovation activity. However, we get very similar estimates if we instead instrument for $[POST_t \times EU8SHR_{a,2004}]$ with $[POST_t \times \widehat{EU8SHR}_{a,2004}]$ where $\widehat{EU8SHR}$ is the vector of predicted values obtained from running OLS of $EU8SHR_{a,2004}$ on a constant, $EU8SHR_{a,1991}$, and $POST_t$ (see Wooldridge [44] or Balli and Sørensen [5]). Throughout, we take the IV results as our preferred estimates of the impact of immigration-induced labor supply shocks on innovation outcomes.

The two key issues with respect to our identification strategy are: what are the sources of variation driving our first stage? And, is this variation unrelated to changes in firm outcomes over the 2004-2008 period (except through immigration inflows)?

Migrant Network Formation

The first question is particularly important given that the countries that joined the EU in 2004 had been behind the Iron Curtain for much of the late twentieth century, so that there was not a continuous flow of migrants from these countries over that period. Nevertheless, strong connections were maintained between pre-Cold War migrants to the UK and their homelands, in part as a response to the adversity that was faced by family members who remained in Eastern Europe. For instance, the Poles – by far the largest migrant group to the UK following the 2004 EU Accession (see Table 4) – formed large communities around the UK as of 1951 (see Holmes [29]), in large part due to a large-scale resettlement plan for Poles following World War II.²⁴ According to the 1951 Census there were 152,000 Polish-born UK residents, constituting the second largest UK immigrant group. These immigrants maintained close connections with family and friends during the Cold War and beginning in 1956 some travel was allowed between Poland and the UK, which facilitated the maintenance of networks. Thus, for members of the Eastern European diaspora the desire to remain connected to their homelands was often heightened during the Cold War.

The fall of the Iron Curtain, between 1989 and 1992, led to migration that was largely driven by these networks, and is the period in which we set our baseline cross-section of EU8 immigrants.²⁵ Future immigration, for instance due to the 2004 EU Accession, was then also likely to be in part

²⁴Known as the Polish Resettlement Act of 1947 it offered British citizenship to Polish troops living in Britain after the war. Most settled in London, Swindon and the industrial areas of the North.

²⁵Note that it would be ideal to use the 1951 distribution of EU8 migrants across UK labor markets – which was almost entirely the product of post-war resettlement – as our baseline cross-section, but these data are not available.

	2004	2005	2006	2007	2008
East of England	15425	22710	23785	23955	18775
East Midlands	9035	18540	23680	22245	16050
London	27860	27920	26605	27925	22255
North East	990	1950	2990	2750	1810
North West	5980	15240	20490	18710	11920
Northern Ireland	3025	8005	8440	8370	5560
Scotland	7180	14445	18035	19630	14180
South East	17320	25570	26950	26105	20315
South West	7275	13875	16975	16430	11930
Wales	2095	5430	6920	6550	3750
West Midlands	6715	16085	19755	18555	13000
Yorkshire and the Humber	5880	15820	19410	17655	12305

Table 5: EU8 Migrant Population by UK Region and Year

driven by these established settlement patterns. And the 1991 immigrant distribution is indeed predictive of the 2004 distribution. We confirm this formally, finding a strong and significant correlation between the 1991 and 2004 immigrant distributions, with an F-Stat of 35, reflecting a persistence in the geography of EU8 migrants over two decades. In Table 4, Row 3 we report the correlation for each of the EU8 immigrant groups, where we see that for most groups the correlation is strong. In fact, the correlation is strongest for Polish immigrants, who subsequently experienced the largest growth in population within the UK following the 2004 Accession. As a result, our first stage is disproportionately driven by this immigrant group and their pre-period settlement patterns. Ultimately, EU8 migrants settled throughout the country post-2004, as Table 5 depicts.²⁶

We note here that, since London is clearly the primary hub for immigration to the UK, Appendix Table A8 presents results of running our baseline empirical analysis with London removed and the findings hold.

Testing the Common Trends Assumption

Networks were therefore very likely to be drivers of the persistence in the immigrant distribution between 1991 and 2004. However, this persistence may also reflect a correlation between economic shocks leading up to 1991 that drove immigration to specific UK locations (or within-UK migration of existing EU8 residents) and that were also correlated with innovation outcomes in the later period. For instance, regions receiving positive economic shocks during the late nineteen eighties may have attracted EU8 immigrants (or internal EU8 migrants), and may have remained desirable locations to live over the subsequent two decades, for reasons that also led to greater firm innovation

²⁶These data come from the UK Worker Registration Scheme which was explicitly set up to monitor the inflows of EU8 migrants into the UK.

post-2004. To address this, we do four things. First, we ask whether employment growth across TTWAs between 1981 and 1991 is correlated with the immigrant distribution in 1991 *as well as* innovation outcomes across TTWAs over the period 2004-2008. Second, we ask whether the share of immigrants in urban areas in 1991 is correlated with the urban share of innovation growth over 2004-2008. And, third, we ask whether a TTWA's distance to London is correlated with both the 1991 share of immigrants and innovation growth over 2004-2008. These three tests help to address the common trends assumption by asking whether there are common correlates of our dependent and independent variables. In other words, we ask whether there are underlying economic trends driving the distribution of EU8 migrants up to 1991, and whether these trends are sufficiently persistent to drive innovation activity over 2004-2008. It is important to note that the existence of correlations would not necessarily indicate that our estimates are biased (e.g., the distribution of immigrants across TTWAs may be dominated by urban areas, but still may be exogenously allocated within that set of areas), but it will help to identify variation that we would be wise to control for in our specifications.

Lastly, we ask whether output per worker and, separately, the average wage in 2004 across TTWAs are correlated with the distribution of EU8 immigrants across TTWAs in 1991. Here we take the 1991 immigrant distribution as given, and simply ask if it is predictive of economic conditions 13 years later. A positive correlation would suggest that the economic factors associated with the location of EU8 immigrants in 1991 may have persisted. We choose these particular outcomes since they are the closest proxies that we have for the serially correlated productivity shocks that we are ultimately concerned about.

Table 6 reports the results of separate OLS regressions in which the five variables discussed above are the dependent variable and the regressors are, separately, the 1991 immigrant distribution and then innovation growth over the 2004-2008 period. We see in column (1) that while there is a positive and significant (at the 5 percent level) correlation between pre-1991 employment growth and the 1991 immigrant distribution, the pre-1991 employment trends are uncorrelated with innovation growth in the more recent period. This suggests that immigrant sorting prior to 1991 was to some extent driven by economic trends, but that these trends did not persist in a manner that impacted later innovation outcomes. On the other hand, column (2) indicates that both the 1991 immigrant distribution and innovation outcomes in the recent period are primarily urban phenomena. Though this does not necessarily indicate a violation of the common trends assumption in our variables of interest, in our baseline regressions we will include a specification that controls for an urban indicator interacted with our treatment. Finally, the last three columns indicate no common correlations with Distance to London, Average 2004 Wage, or 2004 Output per Worker, though we do see that the wage and output levels in 2004 predict subsequent innovation, perhaps unsurprisingly.

Spillovers across TTWAs

With respect to the potential general equilibrium effects discussed in Section 3.4, we first reiterate that it is not clear which direction the effect will go, since out-migration due to falling wages may be offset by in-migration due to output reallocation (the Rybczynski effect). This issue falls within the discussion on the stable unit treatment value assumption (SUTVA) initiated by Rubin [42], and it implies that the estimates should be conservatively interpreted as *relative* treatment effects, rather than absolute effects. For instance, if we assume that low-skill immigration reduces product innovation (recall that it is ambiguous in the model), and if the net effect of this immigration is to induce overall out-migration by existing *high-skill* residents of a TTWA, then our estimates of the product innovation response to the low-skill EU8 immigration shock will be biased away from zero.²⁷ In contrast, if it induces out-migration of low-skill workers then we will understate the magnitude of the effect. At the same time, it is important to note that Dustmann et al. [20] find no out-migration from U.K. regions following immigration supply shocks over the period 1997 to 2005, suggesting that this issue may be unimportant and our estimates may be close estimates of the absolute treatment effects.

5.3 Skill Heterogeneity, Firm Size and a Role for Immigrant Demand

We also explore specific model predictions regarding possible heterogeneity in the extent and direction of the firm innovation response to rising immigrant shares. Formally, we estimate the same series of specifications as above, but now we interact the treatment variable with several pre-period firm-level measures. We estimate:

$$INN_{iat} = c + \lambda_1 [POST_t \times EU8SHR_{a,2004} \times X_{ia}^{init}] + Z_{iat} + \alpha_t + \gamma_i + \epsilon_{iat} \quad (12)$$

where Z_{iat} is a vector of the required two-way interaction terms and X_{ia}^{init} is the value of the interaction term of interest in the initial period, 2004. The use of initial-period values should mitigate the potential endogeneity of these measures.²⁸ We then estimate the reduced form (12) again but replace $EU8SHR_{a,2004}$ with $EU8SHR_{a,1991}$.²⁹ Finally, we estimate an IV version of (12) in which we instrument $[POST_t \times EU8SHR_{a,2004} \times X_{ia}^{init}]$ with $[POST_t \times EU8SHR_{a,1991} \times X_{ia}^{init}]$. Again, we assume that $POST$ and X^{init} are exogenous, and we note that our estimates are very similar when using the predicted EU8 share variable ($\widehat{EU8SHR}$) as above.

We consider each of the main theoretical results. First, we estimate specifications in which product innovation is the dependent variable and in which the treatment is interacted with the (pre-period) *process* innovation indicator. As noted in the theory section, equilibrium requires that both process and product innovation are optimally determined. As a result, the direct (cost-reducing)

²⁷This is because the out-migration of high-skill workers will raise the wage of high-skill workers in that area, making product innovation more costly, thereby reducing its extent.

²⁸Note that the individual terms from the interaction are absorbed in the firm fixed effects.

²⁹Recall that each of these values is multiplied by 100 for ease of interpreting the estimates, as in our baseline specification.

productivity gains associated with process innovation will raise optimal product innovation by raising the profitability of new products. This specification provides evidence on this channel.

We next explore the implications of Proposition 2, which states that a labor supply shock will increase investments in process innovation. We then explore the role of firm size and firm heterogeneity in worker skill in the magnitude of the process innovation response, noting that Lemmas 1 and 2 predict that the response will be increasing in both. As discussed in Section 3.3, Lemma 2 allows us to infer the factor bias of the process innovation being implemented based on the relative magnitude of the response *across* firms of different factor intensities (i.e., we cannot observe the type of process innovation being done, but we can infer this by observing who is doing it). Throughout, we proxy worker skill with the share of employees with a college degree in science or engineering subjects in the pre-period ($SkillShare_{ia}^{pre}$).³⁰ We then proxy firm size with firm revenue in the pre-period ($Revenue_{ia}^{pre}$).³¹

In our final specification we directly explore the demand side impact of the labor supply shock, as summarized in channel 1 of Proposition 1. We first interact the treatment intensity variable with an indicator for whether the firm sold (in the initial period) all of their output locally ($LocalSales_{ia}^{init}$) – defined as within 200 miles of the firm – and, second, an indicator for whether the firm sold all of their output within the UK ($UKSales_{ia}^{init}$). When using this latter indicator we focus on firms who indicate UK sales, but also indicate no sales within 200 miles – i.e., we want to distinguish between the two subsets of firms. To the extent that the local population increase from the EU8 expansion generates greater local demand for goods and services we would expect to observe any product innovation effects locally as well, and to a much lesser (or no) extent UK-wide. We note that a key identifying assumption is that any supply side effects should operate independently of whether firm sales are local versus UK-wide, such that the estimated effects arise from the demand side only.

6 Results

Below we discuss the results of a set of LPM regressions. The results are robust to estimation via conditional logit, which are presented in Appendix D.

6.1 Product Innovation Estimates

Proposition 1 states that the effect of the labor supply shock on product innovation is ambiguous, and depends on the relative strength of the productivity gains associated with EU8 immigrants (arising from the fall in the local average low-skill wage) and the extent of substitution away from product innovation due to its high-skill intensity (due to the rise in the relative high-skill wage) and within-firm product cannibalization. Table 7 presents OLS results based on (11) where the

³⁰Since the relevant “skill” that we are interested in is the skill required to develop and implement new product or process innovations, we believe this measure of science and engineering education is a nearly ideal measure.

³¹In the model firm size is given by firm output, but we lack output prices and so proxy output with revenues.

dependent variable is a binary indicator for whether the firm engaged in product innovation during the 2004-2008 period, and we note again that the pre-treatment period covers 2002-2004. Columns (1)-(4) exploit the 2004 distribution of EU8 immigrant shares while columns (5)-(8) exploit the 1991 distribution. Throughout our empirical analysis we progressively increase the strictness of the specifications, in our final two specifications adding firm fixed effects. In our strictest specifications (columns 4 and 8) we also interact the treatment term with an urban-rural indicator, following our discussion of threats to common trends in Section 5.2. Time fixed effects are included throughout.

We find that it is the substitution effect that dominated, as indicated by the negative and significant coefficient on the treatment variable across all specifications. The magnitude of the coefficient ranges in size from -0.085 to -0.521. Table 9, columns (1)-(3) present IV estimates, where the 1991 distribution instruments for the 2004 distribution. Taking the IV specification with firm fixed effects as preferred (Table 9, column (3)), and applying the coefficient on the treatment variable of -0.085 – and recalling that a one-unit change in the EU8 share is equivalent to a one percentage point change – the estimate indicates that a one percentage point increase in the EU8 immigration share led to an 8.5 percentage point drop in the product innovation rate from the 2004 level. The average observed rise in the EU8 immigrant share across TTWAs due to the EU8 expansion was just over half of one percentage point (see Table 2), which would therefore be associated with a 5.1 percentage point fall in the product innovation rate. On aggregate, product innovation rose three percentage points between 2004 and 2008 (see Table 2), indicating a significant, and countervailing, role for the 2004 immigration shock.

We can also place an upper bound on the economic magnitude by combining these estimates with the innovation premia reported in Table 1. Again, the estimates in Table 1 are simply OLS correlations between the innovation indicators and firm outcomes (with firm fixed effects in some specifications), and are likely biased. However, if we reasonably assume that the Table 1 estimates are upwardly biased, then they can be taken as likely upper bound estimates. If we combine the IV estimates in Table 9, column (3) that causally relate EU8 immigration to product innovation with the Table 1 correlations, we find that the effect of the observed rise in the EU8 immigration share (of just over half of a percentage point) is associated with a $5.1 \times 1.185 = 6$ percent *reduction* in firm revenue, a $5.1 \times 1.001 = 5.1$ percent reduction in employment and a $5.1 \times 1.05 = 5.6$ percent reduction in labor productivity. Again, these effects are upper bounds and, in addition, are offset by increased process innovation, which we explore in the next section.

In Table 10 we explore heterogeneity in the product innovation response based on specification (12). The relevant coefficients are those on the triple interaction terms in which we ask if the product innovation response is increasing or decreasing in worker skill or whether the firm engages in process innovation.³² In columns (1) and (4) we see that the product innovation response is strongest for

³²Note that we do not present theoretical results regarding the role for worker skill in the product innovation response. This is because the role of skill is extremely complex and ultimately ambiguous, as it impacts each of the three channels in Proposition 1, including channel 2 which itself is of ambiguous sign. Nevertheless, we estimate the role for skill since it is a potentially useful empirical result from a policy perspective.

firms that are intensive in low-skill workers – i.e., firms that are relatively intensive in the now-more-abundant factor are even less likely to engage in product innovation relative to other firms following the labor supply shock. Next, in Columns (3) and (6) we interact the treatment with an indicator for whether the firm engaged in process innovation during the pre-period (2002-2004). We note that we restrict this measure to the pre-period in order to isolate variation that is likely to be exogenous to the treatment, so the results should be interpreted as reflecting the impact of product innovation on initially more or less process-intensive firms. The results suggest that the negative product innovation response was mitigated for firms who were engaged in process innovation, consistent with the underlying logic of the model. We note that the results are insignificant when adopting the 1991 EU8 distribution, and only significant at the 10 percent level when using the 2004 distribution. Table 12, column (3) presents IV results that are in line with the OLS estimates, though on the order of a third of the size. Overall, the magnitudes are reasonable and in the direction predicted by the model.

Columns (1)-(4) of Table 13 then ask whether the impact on product innovation is affected by the extent to which the firm sells their output locally (within 200 miles) or, alternatively, within the UK (but not within 200 miles). This specification explores the interaction between EU8 immigrants and local demand for firm output in driving the product innovation decision. This follows from (8), which implies that $\frac{\partial h^*}{\partial M} < 0$. The coefficient on the treatment variable with either measure of local sales interacted indicates that the negative product innovation response is in fact mitigated when the firm sells only locally, and more so when the firm sells within 200 miles relative to UK sales beyond that distance.

We note that these last findings are inconsistent with the product cannibalization channel summarized in Proposition 1, channel 1. As a way of explanation, it may be the case that adding and dropping products takes time, so that our short-run results do not capture long-run declines in product innovation. Or perhaps firms significantly reduce their production of some product lines but do not cease production entirely, which our data would miss. In any case, these results indicate that there is little evidence for a negative *demand-side* effect on product innovation in response to a large-scale immigration shock, suggesting that the observed negative effect is likely driven by channels 2 and/or 3. This result is informative in part because Hottman et al. [30] do find an important role for within-firm product cannibalization in the context of an international trade shock.

Finally, in Appendix Table A8 we present the baseline product innovation results (similar to Table 7) with London removed from the analysis, and we find that the coefficients are somewhat smaller, but otherwise consistent with the baseline results.

6.2 Process Innovation Estimates

Table 8 reports OLS estimates based on the baseline specification (11) in which the dependent variable is now an indicator for process innovation. When the condition in Proposition 2 holds, the model predicts that process innovation will be biased toward the now more abundant factor –

i.e., $\frac{\partial \kappa_L^*}{\partial w_L} < 0$ – such that we expect $\beta_2 > 0$ in (11). This is indeed what we find, with statistically significant coefficients ranging from 0.122 to 0.393. Table 9, columns (4)-(6) present our preferred IV estimates, the strictest of which (column (6)) indicates that firms in areas that had a one percentage point increase in the share of their population from EU8 countries were 11.8 percentage points more likely to be engaged in process innovation after the EU expansion. This is double the actual rise in the average EU8 share, and comes from a baseline mean likelihood of 21 percent (see Table 2) – so is a rather large effect. Again, in Appendix Table A8 we present the baseline results with London removed from the analysis, where we see that the coefficients are consistent with the baseline results.

We can provide additional suggestive evidence on the baseline economic magnitudes by combining these estimates with the innovation premia reported in Table 1. Here we again combine the IV estimates in Table 9, column (6) that causally relate EU8 immigration to process innovation with the Table 1 correlations. We find that the observed rise in the EU8 immigration share over the 2004-2008 period was associated with a $7.1 \times 0.494 = 6$ percent *increase* in firm revenue, a $11.8 \times 0.394 = 5$ percent increase in employment and a $11.8 \times 0.988 = 12$ percent increase in labor productivity.

Table 11, columns (3) and (4), and Table 12, columns (4) and (5), report OLS (Table 11) and IV (Table 12) tests of Lemmas 1 and 2 in the model by introducing two interactions with the treatment variable, the share of workers that are high skilled and the log of firm revenues, as a proxy for firm size. The estimates indicate that process innovation was mitigated by the high-skill content of the firm, such that low-skill-intensive firms increased process innovation relatively more in response to the shock, as evidenced by the negative and significant coefficient. This result is consistent with Lemma 2. In addition, we find that the process innovation response was increasing in firm size. Here we see a positive and significant coefficient, consistent with Lemma 1 – larger firms indeed responded more to the local labor supply shock by increasing process innovation, although the economic magnitude is quite small. In fact, the estimates imply that a one standard deviation increase in firm size increased the impact of EU8 immigrants by about 8 percent relative to the average firm.

7 Concluding Remarks

Immigration policy is a prominent political issue. As a result, it is important to understand the impact that immigrants have on host country economies. The interaction of immigrants and innovation outcomes is an understudied area with potentially large implications for host economy performance in the long run. Here, we explore the product and process innovation responses to labor supply shocks through the lens of a large influx of low-skill immigrants to the UK in 2004. We develop a model and test its predictions, finding that the low-skill labor supply shock 1) increased process innovation while 2) reducing product innovation. We further find evidence in favor of a demand side role for immigrants in spurring new product creation, which was likely offset by a relative increase in the price of high-skill workers, thus discouraging net product innovation activity. The

findings are mostly in line with our model – for instance, both the product and process innovation responses are increasing in the low-skill intensity of firms.

More generally, the results suggest that one reason that the estimated labor market effects of immigration are small is that firms adjust their production processes rapidly in response to changes in input endowments. This is in line with Dustmann and Glitz [21] and their findings that German firms in the tradable sector responded to changes in available skill by altering their production technologies to ones more intensive in the now relatively more abundant skill type. Lewis [36] provided an overview of models that provide channels of adjustment to labor supply shocks, and describes a new thread of research that uses exogenous immigration shocks, such as the one exploited here resulting from the 2004 EU expansion, to test the predictions of this class of models. Our paper suggests that indeed these type of production technology and capital adjustments seem to take place and must be taken into account in any predictions about the impact of immigrants on relative wages of native workers.

The results with respect to product innovation are also informative. Previous work has focused almost exclusively on the direct link between high-skill immigrants and inventive activity. In this paper we explored the implications of a low-skill shock and its direct and indirect effects on product creation. A potential direction for future research is to incorporate both high- and low-skill immigration flows into a model of firm innovation choice– one finding that might differ with our setting is that the demand side effect might be larger in the presence of a larger supply of richer, high-skilled immigrants. This would be particularly relevant for countries for whom immigrant skill is bimodal (such as in the UK or US). A focus on longer run outcomes would also be useful for policy makers.

References

- [1] Acemoglu D. 1998. Why Do New Technologies Complement Skills? Directed Technical Change And Wage Inequality. *The Quarterly Journal of Economics* **113**: 1055–1089.
- [2] Altonji JG, Card D. 1991. The effects of immigration on the labor market outcomes of less-skilled natives. In *Immigration, trade, and the labor market*. University of Chicago Press, 201–234.
- [3] Antras P, Garicano L, Rossi-Hansberg E. 2006. Organizing offshoring: Middle managers and communication costs. Technical report, National Bureau of Economic Research.
- [4] Antràs P, Rossi-Hansberg E. 2009. Organizations and trade. *Annu. Rev. Econ.* **1**: 43–64.
- [5] Balli HO, Sørensen BE. 2013. Interaction effects in econometrics. *Empirical Economics* **45**: 583–603.
- [6] Basker E. 2012. Raising the barcode scanner: Technology and productivity in the retail sector. *American Economic Journal: Applied Economics* **4**: 1–27.
- [7] Bell B, Fasani F, Machin S. 2013. Crime and immigration: Evidence from large immigrant waves. *Review of Economics and statistics* **21**: 1278–1290.
- [8] Blanchflower DG, Lawton H. 2009. The impact of the recent expansion of the eu on the uk labor market. In *EU labor markets after post-enlargement migration*. Springer, 181–215.
- [9] Bloom N, Draca M, Van Reenen J. 2016. Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *The Review of Economic Studies* **83**: 87–117.
- [10] Bloom N, Garicano L, Sadun R, Van Reenen J. 2014. The distinct effects of information technology and communication technology on firm organization. *Management Science* **60**: 2859–2885.
- [11] Bloom N, Sadun R, Van Reenen J. 2012. Americans do it better: Us multinationals and the productivity miracle. *The American Economic Review* **102**: 167–201.
- [12] Card D. 2001. Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics* **19**: 22–64.
- [13] Card D. 2012. Comment: The elusive search for negative wage impacts of immigration. *Journal of the European Economic Association* **10**: 211–215.
- [14] Caselli F, Coleman II WJ, et al. 2006. The world technology frontier. *American Economic Review* **96**: 499–522.
- [15] Cortes P. 2008. The effect of low-skilled immigration on us prices: evidence from cpi data. *Journal of political Economy* **116**: 381–422.

- [16] Dhingra S. 2013. Trading away wide brands for cheap brands. *The American Economic Review* **103**: 2554–2584.
- [17] Duranton G, Puga D. 2001. Nursery cities: Urban diversity, process innovation, and the life cycle of products. *American Economic Review* **91**: 1454–1477.
- [18] Dustmann C, Fabbri F, Preston I, Wadsworth J. 2003. Labour market performance of immigrants in the uk labour market. *Home Office Online Report 05/03* .
- [19] Dustmann C, Frattini T, Halls C. 2010. Assessing the fiscal costs and benefits of a8 migration to the uk*. *Fiscal Studies* **31**: 1–41.
- [20] Dustmann C, Frattini T, Preston IP. 2012. The effect of immigration along the distribution of wages. *Review of Economic Studies* **80**: 145–173.
- [21] Dustmann C, Glitz A. 2015. How do industries and firms respond to changes in local labor supply? *Journal of Labor Economics* **33**: 711–750.
- [22] Eckel C, Neary JP. 2010. Multi-product firms and flexible manufacturing in the global economy. *The Review of Economic Studies* **77**: 188–217.
- [23] Eizenberg A. 2014. Upstream innovation and product variety in the us home pc market. *Review of Economic Studies* **81**: 1003–1045.
- [24] Elsner B. 2013. Emigration and wages: The eu enlargement experiment. *Journal of International Economics* **91**: 154–163.
- [25] Feenstra RC. 1994. New product varieties and the measurement of international prices. *The American Economic Review* : 157–177.
- [26] Fogel M, Peri G. 2016. Immigrants’ effect on native workers: New analysis on longitudinal data. *American Economic Journal: Applied Economics* **8**: 1–34.
- [27] Gaggl P, Wright GC. 2017. A short-run view of what computers do: Evidence from a uk tax incentive. *American Economic Journal: Applied Economics* .
- [28] Hanson GH, Slaughter MJ. 1999. The rybczynski theorem, factor-price equalization, and immigration: evidence from us states. Technical report, National bureau of economic research.
- [29] Holmes C. 2015. *John Bull’s island: immigration and British society, 1871-1971*. Routledge.
- [30] Hottman C, Redding SJ, Weinstein DE. 2014. What is’ firm heterogeneity’in trade models? the role of quality, scope, markups, and cost. *NBER Working Paper Series* : 20436.
- [31] Huergo E, Jaumandreu J. 2004. Firms’ age, process innovation and productivity growth. *International Journal of Industrial Organization* **22**: 541–559.

- [32] Hunt J. 2011. Which immigrants are most innovative and entrepreneurial? distinctions by entry visa. *Journal of Labor Economics* **29**: 417–457.
- [33] Kerr WR. 2013. U.s. high-skilled immigration, innovation, and entrepreneurship: Empirical approaches and evidence. Technical report, National Bureau of Economic Research.
- [34] Lafortune J, Tessada J, Lewis E. 2015. People and machines: A look at the evolving relationship between capital and skill in manufacturing 1860-1930 using immigration shocks. Technical report, National Bureau of Economic Research.
- [35] Lewis E. 2011. Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics* **126**: 1029–1069.
- [36] Lewis E. 2013. Immigration and production technology. *Annu. Rev. Econ.* **5**: 165–191.
- [37] Maré DC, Fabling R, Stillman S. 2014. Innovation and the local workforce. *Papers in Regional Science* **93**: 183–201.
- [38] Markus ML, Robey D. 1988. Information technology and organizational change: causal structure in theory and research. *Management science* **34**: 583–598.
- [39] Ottaviano GI, Peri G. 2012. Rethinking the effect of immigration on wages. *Journal of the European economic association* **10**: 152–197.
- [40] Ozgen C, Nijkamp P, Poot J. 2013. Measuring cultural diversity and its impact on innovation: Longitudinal evidence from dutch firms. *IZA Discussion Paper No. 7129* .
- [41] Peri G, Sparber C. 2009. Task specialization, immigration, and wages. *American Economic Journal: Applied Economics* **1**: 135–169.
- [42] Rubin DB. 1990. Comment: Neyman (1923) and causal inference in experiments and observational studies. *Statistical Science* **5**: 472–480.
- [43] Stuen ET, Mobarak AM, Maskus KE. 2012. Skilled immigration and innovation: evidence from enrolment fluctuations in us doctoral programmes. *The Economic Journal* **122**: 1143–1176.
- [44] Wooldridge J. 1995. M.(2002) econometric analysis of cross section and panel data. *The MIT Press* **0** **5**: 5.

Appendix

A Employment Growth Within and Between UK Firms

In this section we decompose variation in employment between and across UK firms over the period 2004-2010. First, we decompose the change in economy-wide employment into components reflecting the change in employment associated with firms that exist across the entire period; the change in employment due to firm entry (new firms); and the change in employment due to firm exit (firm death). Formally, we calculate:

$$\Delta \sum_i N_i = \Delta \sum_i N_{ip} + \sum_i N_{in} - \sum_i N_{ie}$$

where N is employment and i indexes firms, p indicates “permanent” firms, n indicates “new” firms, and e indicates “exited” firms. This simple calculation produces the following decomposition:

$\Delta \sum N_i$	$\Delta \sum N_{i,p}$	$\sum N_{i,n}$	$\sum N_{i,o}$
2512161	3122572	2670621	3281032

where we see that aggregate employment variation is driven to a large extent by each component. Next, we can further decompose the change in employment associated with permanent firms into between and within components. This produces the following decomposition:

$\Delta \sum N_{i,p}$	Mean	Std. Dev.	Observations
Overall	18.86101	540.799	N=2390584
Between		529.1837	N=1195292
Within		111.4823	T=2

Here we see that, over the period 2004-2010, variation in employment is overwhelmingly explained by differential growth *between* UK firms. For our purposes the key point is simply that the firm growth experience over the period was heterogeneous, suggesting that firms responded quite differently to shocks, a finding that is supportive of the notion that the impact of the EU8 expansion – an important economic event during this period – may have differed significantly across firms.

B Proof of Proposition 2

We want to find the condition for which $\frac{\partial \kappa_L^*}{\partial w_L} < 0$, which is equivalent to $-\frac{\partial F/\partial \kappa_L^*}{\partial F/\partial w_L} < 0$ where the implicit function $F = 0$ is the solution to $\frac{\partial \pi}{\partial \kappa_L} = 0$ (see Section 3.3) and is given by (9). Given the form of the profit function, we can apply the Implicit Function Theorem to re-write this as

$$-\frac{\partial F/\partial \kappa_L^*}{\partial F/\partial w_L} = -\frac{\overbrace{\left[\frac{\partial q}{\partial^2 \kappa_L^*} (p-c) - 2 \frac{\partial q}{\partial \kappa_L^*} \frac{\partial c}{\partial \kappa_L^*} - q \frac{\partial c}{\partial^2 \kappa_L^*} \right]}^1}{\underbrace{\left[\frac{\partial q}{\partial \kappa_L^* \partial w_L} (p-c) - q \frac{\partial c}{\partial \kappa_L^* \partial w_L} - \frac{\partial q}{\partial \kappa_L^*} \frac{\partial c}{\partial w_L} - \frac{\partial q}{\partial w_L} \frac{\partial c}{\partial \kappa_L^*} \right]}_{\substack{4 \\ 5 \quad 6 \quad 7}}} \quad (13)$$

It is clear from (13) that $-\frac{\partial F/\partial \kappa_L^*}{\partial F/\partial w_L} < 0$ and therefore $\frac{\partial \kappa_L^*}{\partial w_L} < 0$ when both the numerator and denominator are either positive or negative (setting aside the negative sign out front). We continue to keep track of the terms and combine partial differentials and find that

$$\frac{\overbrace{[\sigma c(\kappa_L^*)^{\sigma-1} \chi_1 - \sigma + 2] [p - c(\kappa_L^*) + q(\kappa_L^*)]}^{1+3} + \overbrace{2c(\kappa_L^*)^\sigma \chi_2}^2}{\underbrace{w(1 + \kappa_L) [\sigma c(\kappa_L^*)^{(\sigma-1)(1-\sigma)} \chi_3 - \sigma + 1] [p - c(\kappa_L^*) + q(\kappa_L^*)]}_{4+6} - \underbrace{2c(\kappa_L^*)^\sigma \chi_4}_5} > 0 \quad (14)$$

where κ_L^* is the firm's optimal investment in low-skill-biased process innovation and $\chi_1 \equiv \frac{\beta^\sigma w^{1-\sigma} \bar{A}_L^{1+\sigma}}{1+\kappa_L^*}$, $\chi_2 \equiv \frac{\beta^\sigma w^{1-\sigma} \bar{A}_L^{1+\sigma}}{(1+\kappa_L^*)^{\sigma-1}}$, $\chi_3 \equiv \frac{\beta^\sigma w^{1-\sigma} \bar{A}_L^{\sigma-1}}{(1+\kappa_L^*)^{\sigma-1}}$ and $\chi_4 \equiv \frac{\beta^\sigma w^\sigma \bar{A}_L^{\sigma-1}}{(1+\kappa_L^*)^{-\sigma}}$. This is the condition in Proposition 2. ■

C Proof of Lemmas 1 and 2

In Appendix B we found that

$$\frac{\partial \kappa_L^*}{\partial w_L} = \frac{[\sigma c(\kappa_L^*)^{\sigma-1} \chi_1 - \sigma + 2] [p - c(\kappa_L^*) + q(\kappa_L^*)] + 2c(\kappa_L^*)^\sigma \chi_2}{w(1 + \kappa_L) [\sigma c(\kappa_L^*)^{(\sigma-1)(1-\sigma)} \chi_3 - \sigma + 1] [p - c(\kappa_L^*) + q(\kappa_L^*)] - 2c(\kappa_L^*)^\sigma \chi_4} \quad (15)$$

and we want to show that $\frac{\partial \kappa_L^*}{\partial w_L \partial q} < 0$. Given $\sigma > 1$, it is straightforward to show that the numerator is increasing in q , while the denominator is falling in q , such that $\frac{\partial \kappa_L^*}{\partial w_L}$ is increasing in q . The same is true with respect to β so that $\frac{\partial \kappa_L^*}{\partial w_L}$ is also increasing in β – i.e., $\frac{\partial \kappa_L^*}{\partial w_L \partial \beta} < 0$.