

PROCESSING IMMIGRATION SHOCKS: FIRM RESPONSES ON THE INNOVATION MARGIN*

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

September 27, 2017

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, 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 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 we interpret as evidence for 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 our 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 the Secure Data Service of the UK Data Archive for access to the firm-level dataset. Montresor acknowledges the support of an Economic and Social Research Council PhD Studentship.

[†]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 labor supply shocks 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 (2012) 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 (2009)). 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 (2015) and Lewis (2013)) while, possibly at the same time, adjusting their capital stock in response to the labor supply shock (e.g., Lewis (2011) and Ottaviano and Peri (2012)), both of which may mitigate any local wage and employment effects. More generally, Acemoglu (1998) 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. We do this in the context of UK firms' responses to the 2004 expansion of the EU to Eastern Europe.

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. 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 effectively increases the size of the local market for goods and services, leading to the entry of new firms and, thus, more competition among firms. 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.¹ Second, we show that to the

¹See Eckel and Neary (2010), Dhingra (2013) or Hottman et al. (2014).

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. In contrast, the short-run fall in the local low-skill wage due to the increase in labor supply reduces firm production costs, which raises the profitability of all products, and this leads firms to *increase* 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.

There is little extant literature relating labor supply shocks to process innovations.² At the same time, organizational changes (process innovations) have been shown to be a key aspect of the firm response to technology adoption³ and international trade,⁴ suggesting that labor supply shocks may also induce these types of responses. Here we provide evidence to this effect. Furthermore, these effects may be economically important; for 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.⁵ Finally, we are unaware of existing evidence on the role that labor supply shocks play in firms' decisions regarding optimal product mix and scope (product innovations), and again we provide evidence on this.

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. (2012) 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 (2011) 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 (2013) 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. (2013)). Our model also assumes that product innovation requires high-skill labor as the main input but our focus throughout the paper

²The closest paper to ours is Maré et al. (2014) who study the relationship between innovation and immigration in New Zealand, but do not employ an identification strategy.

³See Markus and Robey (1988) for an early discussion and Bloom et al. (2014), Bloom et al. (2012) and Gaggl and Wright (2017) more recently.

⁴See, e.g., Antras et al. (2006) and Antràs and Rossi-Hansberg (2009).

⁵Financial Times article November 2, 2016: <https://www.ft.com/content/e514de74-a0e3-11e6-86d5-4e36b35c3550>.

is the impact of a low-skill labor supply shock.

After presenting our model, we bring its predictions to UK data by exploiting the expansion of the European Union (EU) in 2004 as a differential shock to the supply of low-skill labor across UK local labor markets. Using firm-level panel data on product and process innovation activities, we estimate several specifications that produce consistent results. We first find that this large, low-skill labor supply 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. To this point, in Appendix A we show that 77 percent of the variation in total UK employment growth over the period 2004-2010 occurred *between* firms, suggesting that firms may not only be subject to heterogeneous shocks but also that firms may respond to the same shocks in heterogeneous ways. First, we find that the response was increasing in firms' low-skill production intensity as well as firm size. And second, we find that product innovation fell in response to the migration, though less so within low-skill intensive firms and less so 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 finding.

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

We bring the predictions of the model to the data by exploiting a large shock to the relative supply of low-skill labor across UK travel to work areas (TTWAs are standardized UK geographic units, described further in Section 4 below) in the form of the expansion of the EU in 2004. The expansion brought in eight Central and Eastern European countries: Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia and Slovenia. Though citizens of these countries were immediately granted 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 (2009) 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

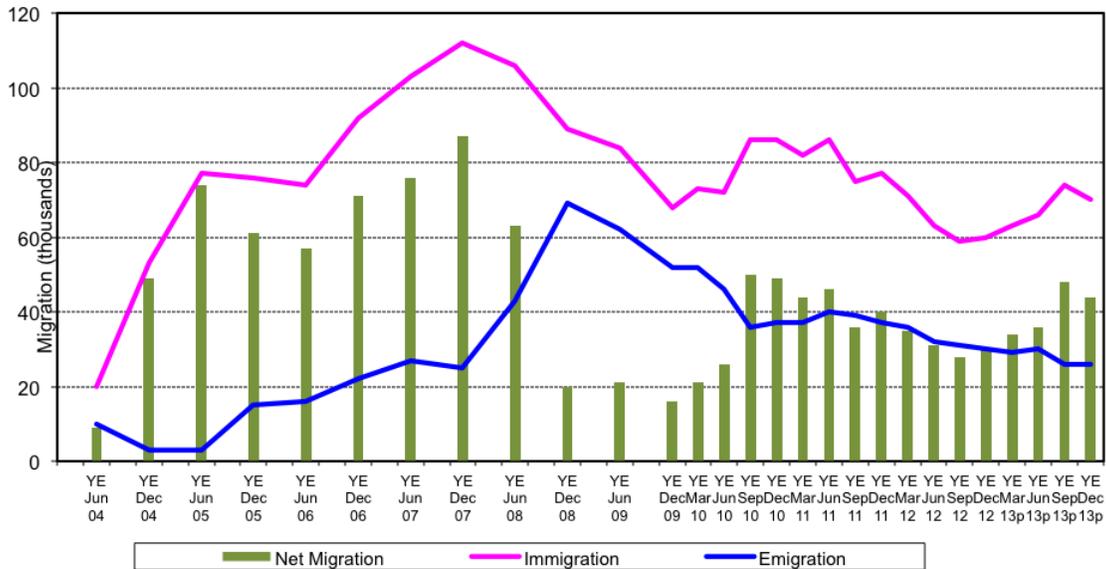


Figure 1: EU8 Immigration, 2004-2013. Source: ONS

EU8 countries. Figure 2 depicts the long-run trend in immigration to the UK, indicating that 2004 represented a significant departure from trend. 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 (2013)).

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 new countries to the UK would be 5,000-13,000.⁶ 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

⁶These figures were generated by Dustmann et al. (2003).

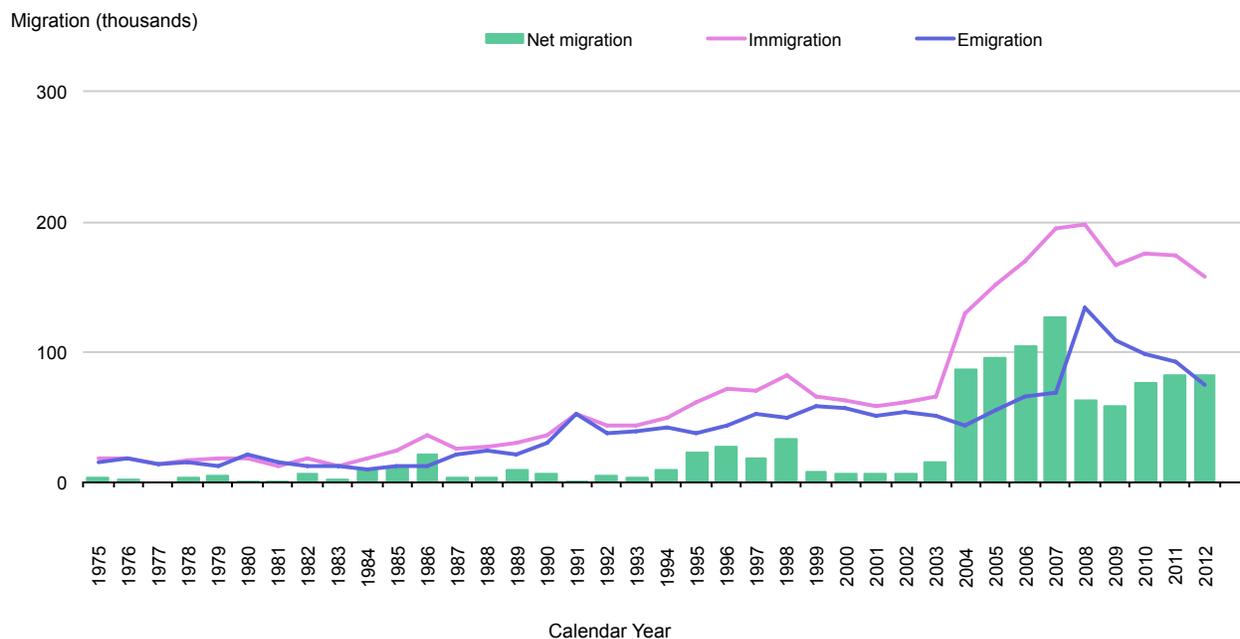


Figure 2: Long-Run Trend in UK Immigration, 1975-2012. Source: ONS

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.⁷ According to Dustmann et al. (2010) 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 (2009) 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 low-skill natives and immigrants are imperfectly substitutable, this fall in the average low-skill wage would 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)

⁷This 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. (2010)).

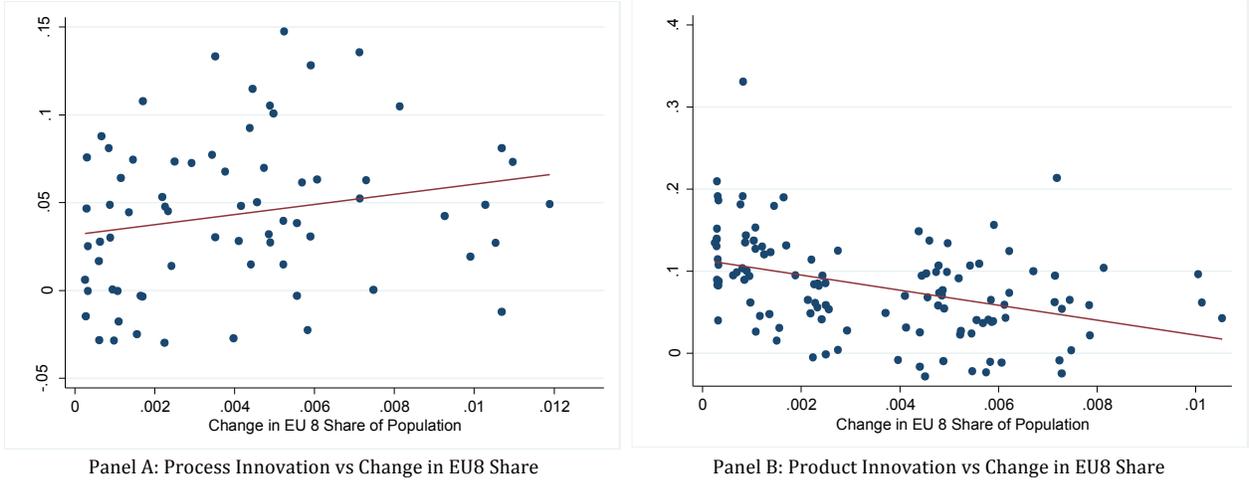


Figure 3: Change in Innovation vs Change in EU8 Immigrants Across UK Travel to Work Areas, 2004-08.

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 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 an identification strategy to explore these relationships in more detail.

3 Model

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:

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

where q_0 represents consumption of the homogeneous good, $q_j^m \equiv \int_i q^m di$ is the agent's consumption of brand j varieties, $Q^m \equiv \int q_j^m dj$ 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 $q_0^m + \int \int_i p q^m di dj = I^m$, 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 :

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

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. (2014) 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 (2013) and one that Hottman et al. (2014) 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 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.⁸

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.

⁸An 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.

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 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 $r_l \kappa_l$.⁹ Here we note that the

⁹In 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. (2006) 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

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 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

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

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.

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.¹⁰

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 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 (2013), 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. First, plugging in for the optimal quantity, from (7), and the optimal profits – noting that optimal process innovation must also be incorporated, and is solved for below – and differentiating (8) with respect to the low-skill wage, we find that $\frac{\partial h^*}{\partial w_L} < 0$. By reducing production costs, the low-skill labor supply shock makes production of all varieties more profitable, which increases the equilibrium range of profitable varieties.

Second, differentiating (8) with respect to M , the size of the local market, we find that $\frac{\partial h^*}{\partial M} < 0$.¹¹ This is the somewhat counter-intuitive result that is analogous to the trade context described in Eckel and Neary (2010) and Dhingra (2013). In short, firms respond to the overall rise in

¹⁰Of course, the cost function explicitly incorporates the firm’s optimal choice of factors.

¹¹Note that q^* and π^* are functions of M .

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.¹² 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.

Finally, since low-skill labor and high-skill labor are imperfect substitutes, the fall in the low-skill wage leads to an increase in the high-skill wage. 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. $\frac{\partial h^*}{\partial w_L} < 0$. *By reducing production costs, a low-skill labor supply shock increases the range of profitable varieties, thereby **increasing product scope**.*
2. $\frac{\partial h^*}{\partial M} < 0$. *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**.*
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**.*

Proposition 1 indicates that the direction of the product innovation response to a low-skill labor supply shock ultimately depends on the relative values of model parameters, and so is an empirical question. This is because the relative increase in the supply of low-skill labor generates productivity gains for the firm that increase product innovation (channel 1), but also increases the fixed costs associated with product innovation, thereby reducing product innovation (channel 3). At the same time, the market for all products is now more competitive and this is also a force for reducing product innovation (channel 2).

The condition for optimal product scope, (8), also implies heterogeneity in the extent of the response to a low-skill local labor supply shock. This is formalized in the following Lemma.

¹²In Eckel and Neary (2010) and Dhingra (2013) the expansion in market size is due to international trade, rather than migration.

Lemma 1 (Role of Factor Intensity in Product Innovation). *The product innovation response to a low-skill local labor supply shock is more positive (or less negative) the larger the firm's initial intensity in low-skill labor.*

This is because the effect from channel 1 is increasing in the low-skill intensity of the firm, while the effect due to channel 3 is falling in the firm's low-skill intensity. Channel 2 is unaffected by the skill intensity of the firm. The combined impact from channels 1 and 2 works toward increasing product scope, or possibly reducing product scope less than it would otherwise be.

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 innovation 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.¹³

¹³The complexity of the condition arises due to the endogenous responses of product innovation and output.

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.¹⁴ We formalize this and a further implication in the following lemma:

Lemma 2 (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 3 (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 2 and 3 can be found in Appendix C.

Lemma 3 is 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 factor bias of the process innovation response will be in the direction of whatever factor is now more abundant. Furthermore, when this is true it will also be the case that 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 some 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 equilibrium¹⁵ are that both labor markets and output markets must

¹⁴Note 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.

¹⁵Of course, depending on the complexity of the model there may be many general equilibrium consequences that we abstract from here.

clear at the national level (assuming no international trade) in general equilibrium. 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 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 (1999)).

Thus, there may be spillovers across areas that we abstract from in the partial equilibrium described above. We are careful to address the implications of these general equilibrium effects 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 UK Travel to Work Areas (TTWAs), a geographic unit developed by the Office of National Statistics (ONS) precisely for the purpose of identifying local labor markets.¹⁶ In short, these labor markets are defined in order to cover both metropolitan areas as well as their commuter suburbs.¹⁷ 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.

In the main empirical section we exploit cross-sectional variation in EU8 immigrant shares from the 1991 Census, which we use to predict subsequent inflows over the 2004-2008 period, as we discuss further below. We then combine the 1991 Census data with firm-level panel data on innovation activities from three waves of the Community Innovation Survey (CIS), covering the period 2002 to 2008. The CIS is the 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 in our case we exploit survey responses regarding firms’ innovation activities between 2002 and 2004 – the period (mostly) prior to the EU8 accession – as well as between 2004 and 2006 and 2006 and 2008. Table 2 presents basic descriptive information

¹⁶We use the 2001 ONS definition of a TTWA.

¹⁷Formally, 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”.

on the CIS variables. 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.¹⁸ 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, the EU enlargement occurred on May 1st 2004, whereas we associate 2004 with our pre-period for variables related to firm outcomes (we do not rely on 2004 variation in immigrant inflows, our right hand side variable). 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, and 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.

In some specifications we control for pre-period trends, though we cannot do this using the CIS which only goes back to 2002. We therefore exploit firm-level data, which we aggregate to the TTWA level, from the Annual Respondents Database (ARD). The ARD is a census of large firms and a repeated sample of smaller ones, constituting around 60-75,000 firms per year. For a comprehensive description of this dataset see Criscuolo et al. (2003) or for a summary see Breinlich and Criscuolo (2011).

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.¹⁹ 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

¹⁸We 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.

¹⁹This is paraphrased from the 2008 CIS.

not want to rely on R&D expenditure entirely because not all expenditures will successfully lead to implementation of new products or processes.

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 (2012), 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 1 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. 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. 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.

Variable	Coefficient	Standard Error
Improve Product Quality	0.018*	0.010
Improve Production Flexibility	0.055***	0.010
Improve Production Capacity	0.073***	0.009
Reduce Per Unit Costs	0.061***	0.008
Improve Health and Safety	0.004	0.007
Increase Value Added	0.011	0.008
Capital Acquisition (millions £)	0.059	0.191

Table 1: Correlates with Process Innovation

5 Specifications & Identification

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,

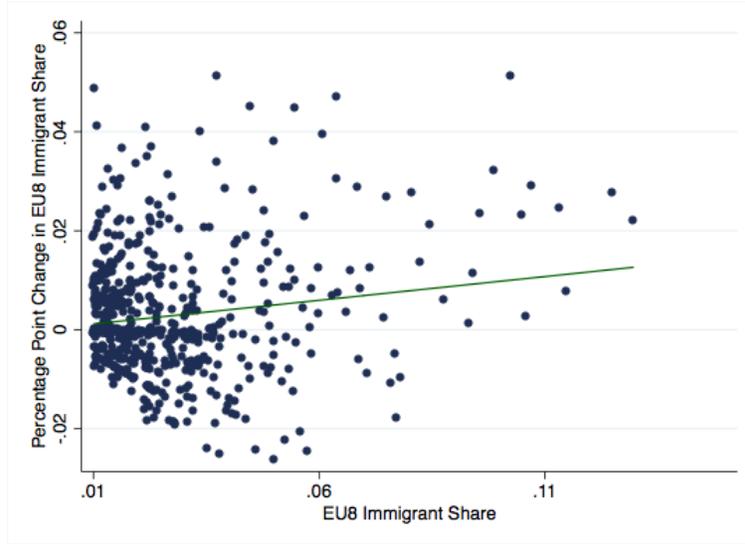


Figure 4: EU8 TTWA Immigrant Share, 2004 vs Percentage Point Change in Share, 2004-2008

we begin by considering the following “naive” difference-in-differences specification (note that we do not estimate this):

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

where INN is one of the binary innovation measures of interest, associated with firm i located in TTWA a in period t ; $EU8SHR$ is the share of EU8 immigrants in TTWA a in year 2004; $POST$ is an indicator equal to 1 for post-2004 periods and 0 for the 2002-2004 period; and α_t and γ_i are period and firm fixed effects, respectively. Here, 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 (1991) and Card (2001). 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 immi-

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

Nevertheless, 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) using a lagged EU8 share variable, $EU8SHR_{a,1991}$, reflecting the share of EU8 immigrants in a TTWA in 1991.²⁰ 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 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 test in the next section.

But once again we are relying on the strength of the ethnic enclave approach, and in going back to 1991 we are relying on its persistence over two decades. Reassuringly, 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. Again, in the next section we go further to confirm that the persistence is due to the enclave-driven channel, rather than a persistence in productivity shocks over time.

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 (1990), 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.²¹ In contrast, if it induces out-migration of low-skill workers then we will understate the

²⁰Bell et al. (2013) use the 2001 Census distribution of immigrants from EU8 countries to implement a similar procedure and isolate the exogenous component of immigrants across locations within the UK.

²¹This is because the out-migration of high-skill workers will raise the wage of high-skill workers in that area, making

magnitude of the effect. At the same time, it is important to note that Dustmann et al. (2012) 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.

We estimate the following specification:

$$INN_{iat} = c + \beta_1 [POST_t \times EU8SHR_{a,1991}] + \beta_2 X_{iat} + \beta_3 \Delta WAGE_a^{Pre} + \beta_4 \Delta VAW_a^{Pre} + \alpha_t + \gamma_i + \epsilon_{iat} \quad (12)$$

where X_{iat} is a set of firm-level control variables and in our strictest specifications we also control for pre-period trends by including the change in the average wage in a TTWA ($WAGE_a^{Pre}$) and value-added per worker (VAW_a^{Pre}) in a TTWA over the period 1997-2003, variables drawn from the ARD (see Section 4 above). We again note that we do not have innovation measures prior to 2002, so we draw from the ARD in order to control for trends in variables that are likely to be highly correlated with innovation, namely wages and value added. We estimate (12) with both a linear probability model (LPM) as well as a conditional logit (CL) regression. Since, in this specification, the right-hand-side variable of interest 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. In Section 6 below we discuss these estimates.

As a robustness check, in an online appendix we also report estimates in which we interact the 2004 dummy with the *predicted* distribution of immigrants in 2004. In other words, we construct a new variable to proxy for the underlying treatment effect. Specifically, the predicted distribution is constructed as a shift-share measure, much like the measure in Altonji and Card (1991). First, we construct a proxy for the 2004 immigrant distribution across TTWAs by allocating all U.K. inflows of immigrants from a particular country between 1991 and 2004 to TTWAs based on their 1991 distribution. We then sum across immigrant origin countries to obtain a TTWA-specific measure of predicted immigrant shares for 2004. We find very similar results using this measure.

5.2 A Test of the Identification Strategy

In the context of our research design, the estimates will be unbiased as long as the 1991 distribution of EU8 immigrants across TTWAs impacts firm innovation outcomes only through its enclave-driven effect on the observed 2004 EU8 distribution (which strongly predicts post-2004 EU8 inflows). In other words, the 1991 EU8 distribution must not also be correlated with other determinants of innovation outcomes. As a test of this assumption, we ask whether the change over 2002 to 2008 in output per worker and, separately, the change in the average wage across TTWAs are correlated

product innovation more costly, thereby reducing its extent.

with the distribution of EU8 immigrants across TTWAs in 1991. We choose these outcomes since they are the closest proxies that we have for the potential productivity shocks that we are concerned about.

To do this, we estimate specification (12) but now our dependent variable is either the change over 2002 to 2008 in output per worker or, separately, the change in the average wage. In short, these specifications produce insignificant coefficients for β_1 (p-values of 0.41 and 0.36, respectively), and we therefore conclude that the correlation between the 1991 immigrant distribution and the 2004 distribution (and therefore subsequent changes in the immigrant distribution over 2004-2008) is likely to be driven by enclave-based network effects, rather than serially correlated productivity shocks over the two decades.

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. First, Lemma 1 indicates that the product innovation response will be larger, or less negative, for firms that are intensive in low-skill labor. Second, Lemma 2 indicates that the firm process innovation response should be increasing in firm size. Third, Lemma 3 indicates that the process innovation response should be rising in the low-skill intensity of the firm. As discussed in Section 3.3, this result 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). Fourth, we can explore the relevance of channel (2) from Proposition 1, which predicts a role for the demand side variation generated by immigrant inflows. And, fifth, we explore the effect that process innovations have on the extent of product innovation since in the model they are jointly determined, and complementary.

Formally, we estimate a second set of specifications in which we interact the treatment variable with several firm-level measures. We estimate:

$$\begin{aligned} INN_{iat} = c + \lambda_1 [POST_t \times EUSSHR_{a,1991} \times X_{ia}^{pre}] \\ + \lambda_2 \Delta WAGE_a^{Pre} + \lambda_3 \Delta VAW_a^{Pre} + Z_{iat} + \alpha_t + \gamma_i + \epsilon_{iat} \end{aligned} \quad (13)$$

where Z_{iat} is a vector of the required two-way interaction terms and X_{ia}^{pre} is the value of the interaction term of interest in the pre-period, 2002-2004. The use of pre-period values should mitigate the potential endogeneity of these measures.²²

We consider each of the main theoretical results. First, Proposition 1 implies that the role of skill in the product innovation response to a labor supply shock is key. This is because channels

²²Note that the individual terms from the interaction are absorbed in the firm fixed effects.

(1) and (2) affect product innovation in the same direction as outlined in the Proposition regardless of the skill content of the inflows, but channel (3) produces a negative effect only in the case when immigration is relatively low-skilled. Thus, under a high-skill labor supply shock the product innovation response is unambiguously positive, whereas it is ambiguous in the case of a low-skill shock, the case captured by the Proposition and that we explore in the empirics.

As an additional test of the model, we also estimate a specification 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) 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. Throughout, we proxy worker skill with the share of employees with a college degree in science or engineering subjects in the pre-period ($SkillShare_{ia}^{pre}$).²³ We then proxy firm size with firm revenue (turnover) in the pre-period ($Turnover_{ia}^{pre}$).²⁴

In our final specification we directly explore the demand side impact of the labor supply shock, as summarized in channel (2) of Proposition 1. We first interact the treatment intensity variable with an indicator for whether the firm sold (in the pre-period) all of their output locally ($LocalSales_{ia}^{pre}$) – defined as within 200 miles of the firm – and, second, an indicator for whether the firm sold all of their output within the UK ($UKSales_{ia}^{pre}$). 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.

²³Since 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.

²⁴In the model firm size is given by firm output, but we lack output prices and so proxy output with revenues.

6 Results

Below we discuss the results of a set of LPM regressions. Importantly, the results are robust to using the conditional logit model and to the inclusion of controls for labor productivity, results that we show in Appendix D. In addition, here we adopt the 1991 distribution of EU8 immigrants across TTWAs in constructing our treatment variable (interacted with a post-2004 dummy) though, as discussed above, we also estimated specifications using the “predicted” 2004 distribution. The results are qualitatively similar when using this shift-share approach and these results are also available 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. In Table 3 the dependent variable is a binary indicator for whether the firm engaged in product innovation during the 2004-2008 period, where we note again that the pre-treatment period covers 2002-2004 and the pre-trend controls cover 1997-2003. Here we find that it is the substitution effect that dominates, as indicated by the negative coefficient on the treatment variable. The magnitude of the coefficient ranges in size from $-.193$ to $-.521$, with a significant reduction in the effect due to the addition of the control variables, but little change due to the addition of firm fixed effects suggesting that within-firm (rather than across-firm) variation is relatively important. Taking the specification with firm fixed effects as preferred (column (4)), and applying the coefficient on the treatment variable of $-.193$, the estimate indicates that a one percentage point increase in the EU8 share (which is in line with the observed rise due to the EU8 expansion, see Figure 3) led to a 7.6 percentage point drop in the product innovation rate from the 2004 level. In fact, on aggregate, product innovation dropped 9 percentage points between 2004 and 2008, suggesting a significant role for the 2004 immigration shock.

In Table 5, columns (3) through (4), we explore heterogeneity in the product innovation response based on specification (13). 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. In column (3) we see that the product innovation response is strongest for firms that are intensive in low-skill workers – i.e., consistent with the model, 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. In Column (4) 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. The results suggest that the negative product innovation response is mitigated for firms who were engaged in process innovation, consistent with the underlying logic of the model.

Columns (3) and (4) of Table 6 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). As discussed, this specification explores the role for local demand for a firm’s output in driving the product innovation decision. 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 characterized by Proposition 1, channel (2), though they should be taken as merely suggestive since the mechanism is clearly not cleanly identified. 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. This result is informative in part because Hottman et al. (2014) do find an important role for within-firm product cannibalization in the context of an international trade shock.

6.2 Process Innovation Estimates

Column (1) of Table 4 reports estimates based on the baseline specification (12) 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 (12). This is indeed what we find, with a statistically significant coefficient of 0.165, indicating that firms in areas that had a one percentage point increase in the share of their population from EU8 countries were 16.5 percent more likely to be engaged in process innovation after the EU expansion. This result is robust to including controls for firm sales or the share of workers with a college degree in science or engineering fields as well as the inclusion of firm or industry fixed effects. Again the pattern of results suggests that the response is driven by within-firm variation.

Table 5, columns (1) and (2) report tests of Lemmas 2 and 3 in the model by introducing two interactions with the treatment variable, the share of workers that are high skilled and the log of firm sales, as a proxy for firm size. Column (1) indicates that process innovation is mitigated

by the high-skill content of the firm, such that low-skill-intensive firms increase process innovation relatively more in response to the shock, as evidenced by the negative and significant coefficient. This result is consistent with Lemma 3. Next, column (2) explores whether the process innovation response is increasing in firm size. Here we see a positive and significant coefficient, consistent with Lemma 2 – larger firms indeed respond more to the local labor supply shock by increasing process innovation, although the economic magnitude is small.

Table 6 reports estimates of the extent to which the process innovation response differed depending on the location of the target market for the firm’s output. In this case it is especially important to be cautious in interpreting the results since there is no direct theoretical mechanism. However, to the extent that process and product innovation are complementary, and to the extent that the immigration-induced demand shock affects product innovation, we may expect to see a response on the process innovation margin as well. In Table 6 we see that firms that are more likely to be impacted on the demand side due to the EU8 expansion in fact are more likely to engage in process innovation, consistent with the indirect mechanism implied by the model.

7 Concluding Remarks

Various countries, including the UK, are considering tightening their immigration policies. 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, in this case, was overwhelmed by the fact that the relative price of high-skill workers increased, thus discouraging net product innovation activity. Our findings are in line with our model, as further evidenced by the fact that the process innovation response is shown to be increasing in the low-skill intensity of firms while the product innovation response is decreasing in low-skill intensity.

More generally, this paper suggests that one reason why the estimated labor market effects of immigration are so small is that firms adjust their production processes fairly rapidly to changes in input endowments. 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. This would be particularly

relevant for any country where immigrant skill is bimodal (such as in the UK or US).

References

- Acemoglu D. 1998. Why Do New Technologies Complement Skills? Directed Technical Change And Wage Inequality. *The Quarterly Journal of Economics* **113**: 1055–1089.
URL <http://ideas.repec.org/a/tpr/qjecon/v113y1998i4p1055-1089.html>
- 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.
- Antras P, Garicano L, Rossi-Hansberg E. 2006. Organizing offshoring: Middle managers and communication costs. Technical report, National Bureau of Economic Research.
- Antràs P, Rossi-Hansberg E. 2009. Organizations and trade. *Annu. Rev. Econ.* **1**: 43–64.
- Basker E. 2012. Raising the barcode scanner: Technology and productivity in the retail sector. *American Economic Journal: Applied Economics* **4**: 1–27.
- Bell B, Fasani F, Machin S. 2013. Crime and immigration: Evidence from large immigrant waves. *Review of Economics and statistics* **21**: 1278–1290.
- 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.
- 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.
- 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.
- Breinlich H, Criscuolo C. 2011. International trade in services: A portrait of importers and exporters. *Journal of International Economics* **84**: 188–206.
- Card D. 2001. Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics* **19**: 22–64.
- Card D. 2012. Comment: The elusive search for negative wage impacts of immigration. *Journal of the European Economic Association* **10**: 211–215.
- Caselli F, Coleman II WJ, et al. 2006. The world technology frontier. *American Economic Review* **96**: 499–522.
- Criscuolo C, Haskel J, Martin R. 2003. Building the evidence base for productivity policy using business data linking. *Economic Trends* **600**: 39–51.

- Dhingra S. 2013. Trading away wide brands for cheap brands. *The American Economic Review* **103**: 2554–2584.
- Dustmann C, Fabbri F, Preston I, Wadsworth J. 2003. Labour market performance of immigrants in the uk labour market .
- Dustmann C, Frattini T, Halls C. 2010. Assessing the fiscal costs and benefits of a8 migration to the uk*. *Fiscal Studies* **31**: 1–41.
- Dustmann C, Frattini T, Preston IP. 2012. The effect of immigration along the distribution of wages. *Review of Economic Studies* **80**: 145–173.
- Dustmann C, Glitz A. 2015. How do industries and firms respond to changes in local labor supply? *Journal of Labor Economics* **33**: 711–750.
- Eckel C, Neary JP. 2010. Multi-product firms and flexible manufacturing in the global economy. *The Review of Economic Studies* **77**: 188–217.
- Elsner B. 2013. Emigration and wages: The eu enlargement experiment. *Journal of International Economics* **91**: 154–163.
- Gaggl P, Wright GC. 2017. A short-run view of what computers do: Evidence from a uk tax incentive. *American Economic Journal: Applied Economics* .
- Hanson GH, Slaughter MJ. 1999. The rybczynski theorem, factor-price equalization, and immigration: evidence from us states. Technical report, National bureau of economic research.
- 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.
- Hunt J. 2011. Which immigrants are most innovative and entrepreneurial? distinctions by entry visa. *Journal of Labor Economics* **29**: 417–457.
- Kerr WR. 2013. Us high-skilled immigration, innovation, and entrepreneurship: Empirical approaches and evidence. Technical report, National Bureau of Economic Research.
- Lewis E. 2011. Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics* **126**: 1029–1069.
- Lewis E. 2013. Immigration and production technology. *Annu. Rev. Econ.* **5**: 165–191.
- Maré DC, Fabling R, Stillman S. 2014. Innovation and the local workforce. *Papers in Regional Science* **93**: 183–201.

- Markus ML, Robey D. 1988. Information technology and organizational change: causal structure in theory and research. *Management science* **34**: 583–598.
- Ottaviano GI, Peri G. 2012. Rethinking the effect of immigration on wages. *Journal of the European economic association* **10**: 152–197.
- Ozgen C, Nijkamp P, Poot J. 2013. Measuring cultural diversity and its impact on innovation: Longitudinal evidence from dutch firms .
- Peri G, Sparber C. 2009. Task specialization, immigration, and wages. *American Economic Journal: Applied Economics* **1**: 135–169.
- Rubin DB. 1990. Comment: Neyman (1923) and causal inference in experiments and observational studies. *Statistical Science* **5**: 472–480.
- 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.

Table 2: Summary Statistics, 2004 and 2008

2004								
	Process Inn	Product Inn	EU-8 Share	High-Skill Share	Log Sales	Log Productivity	Share Selling w/in 200 miles	Share Selling in UK
N	8555	8555	8555	8555	8555	8555	8555	8555
Mean	0.21	0.15	0.3	0.14	1.62	-2.7	0.29	0.22
SD	0.41	0.34	0.47	0.89	1.89	1.09	0.45	0.22
2008								
	Process Inn	Product Inn	EU-8 Share	High-Skill Share	Log Sales	Log Productivity	Share Selling w/in 200 miles	Share Selling in UK
N	8555	8555	8555	8555	8555	8555	8555	8555
Mean	0.28	0.17	1.39	0.09	1.89	-2.51	0.13	0.16
SD	0.45	0.38	1.34	0.53	1.91	1.04	0.34	0.36

Notes: The table presents summary statistics across UK Travel to Work Areas for variables used in the empirical specifications. The variables come from the UK Community Innovation Survey, 2004 and 2008, except for the EU-8 Share of immigrant workers, which comes from the UK Quarterly Labour Force Survey.

Table 3: Baseline Product Innovation Effects

	(1)	(2)	(3)	(4)	(5)
	Product Innovation	Product Innovation	Product Innovation	Product Innovation	Product Innovation
EU8 Share x Post-2004	-0.521*** (-3.62)	-0.227*** (-4.84)	-0.193** (-2.81)	-0.144** (-2.16)	-0.170*** (-4.79)
EU8 Share	11.29 (1.08)	19.06 (1.71)			20.01 (1.60)
Log Turnover		-0.581** (-3.11)	-0.176 (-0.76)	-0.551 (-0.18)	-0.596** (-3.15)
Skill Share		0.506 (1.50)	0.647 (1.89)	0.390 (0.67)	0.504 (1.51)
Wage Change				0.036 (0.67)	-0.006 (-0.43)
Output Change				-0.003 (-1.14)	0.001 (1.11)
Observations	8555	8089	8089	8078	8078
Year FE	Yes	Yes	Yes	Yes	Yes
Other FE	No	No	Firm	Firm	Industry

Notes: EU8share is the share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We control for firm-level time-varying turnover (sales) and the skill share of employment. We estimate a linear probability model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Standard errors are clustered at the TTWA level. t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0$.

Table 4: Baseline Process Innovation Effects

	(1)	(2)	(3)	(4)	(5)
	Process Innovation	Process Innovation	Process Innovation	Process Innovation	Process Innovation
EU8 Share x Post-2004	0.165* (2.06)	0.143** (2.32)	0.135** (2.04)	0.113** (2.35)	0.144** (2.35)
EU8 Share	-0.479 (-1.93)	-0.958*** (-3.94)			-0.760** (-2.65)
Log Turnover		0.035*** (7.30)	0.042 (1.56)	0.043 (1.79)	0.034*** (7.12)
Skill Share		0.0172 (1.63)	-0.042* (-2.23)	-0.042* (-2.23)	0.015 (1.37)
Wage Change				-0.001 (0.71)	-0.00 (-0.46)
Output Change				0.00 (0.78)	0.0001*** (3.47)
Observations	8555	8089	8089	8078	8078
Year FE	Yes	Yes	Yes	Yes	Yes
Other FE	No	No	Firm	Firm	Industry

Notes: EU8 Share is the share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We control for firm-level time-varying turnover (sales) and the skill share of employment. We estimate a linear probability model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.001.

Table 5: Interactions Models

	(1)	(2)	(3)	(4)
	Process Innovation	Process Innovation	Product Innovation	Product Innovation
EU8 Share x Post-2004	0.969* (2.30)	1.217* (2.07)	1.259*** (3.39)	0.746* (2.20)
EU8 Share x Post-2004 x Skill Share	-5.087** (-2.50)		-3.022*** (-0.01)	
EU8 Share x Post-2004 x Log Turnover		0.001*** (15.32)		
EU8 Share x Post-2004 x Process Innovation				1.192* (2.11)
Wage Change	-0.001 (-0.71)	-0.001 (-0.71)	0.001 (0.42)	0.001 (0.52)
Output Change	0.00 (0.84)	0.00 (0.84)	0.00 (-1.18)	0.00 (-1.42)
Observations	8078	8078	8078	8078
Year FE	Yes	Yes	Yes	Yes
Other FE	Firm	Firm	Firm	Firm

Notes: EU8share is the share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a linear probability model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Note that the individual controls from the interaction terms are absorbed by the firm fixed effects (e.g., the pre-period skill share). The two-way interaction terms are not absorbed but for brevity we do not report them here. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.001.

Table 6: Demand Side Effect

	(1)	(2)	(3)	(4)
	Process Innovation	Process Innovation	Product Innovation	Product Innovation
EU8 Share x Post-2004	0.692 (1.55)	0.938* (2.28)	-0.005 (-0.08)	-0.005 (-0.09)
EU8 Share x Post-2004 x Sold Locally	3.324*** (-6.37)		1.575*** (-4.03)	
EU8 Share x Post-2004 x Sold Within UK		1.686*** (-4.58)		1.354* (-2.05)
Wage Change	-0.001 (-0.86)	-0.001 (-0.86)	0.001 (0.41)	0.001 (0.41)
Output Change	0.00 (0.83)	0.00 (0.92)	0.00 (-1.62)	0.00 (-1.57)
Observations	8524	8524	8544	8544
Year FE	Yes	Yes	Yes	Yes
Other FE	Firm	Firm	Firm	Firm

Notes: EU8share is the share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a linear probability model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Sold Locally is an indicator variable for whether all firm output was sold within 200 miles of production. Note that the individual controls from the interaction terms are absorbed by the firm fixed effects (e.g., the pre-period skill share). The two-way interaction terms are not absorbed but for brevity we do not report them here. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.001.

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

$$-\frac{\partial F/\partial \kappa_L^*}{\partial F/\partial w_L} = -\frac{\overbrace{\left[\frac{\partial q}{\partial^2 \kappa_L^*} (p-c) - 2 \frac{\partial q}{\partial \kappa_L^*} \frac{\partial c}{\partial \kappa_L^*} - q \frac{\partial c}{\partial^2 \kappa_L^*} \right]}^1}{\underbrace{\left[\frac{\partial q}{\partial \kappa_L^* \partial w_L} (p-c) - q \frac{\partial c}{\partial \kappa_L^* \partial w_L} - \frac{\partial q}{\partial \kappa_L^*} \frac{\partial c}{\partial w_L} - \frac{\partial q}{\partial w_L} \frac{\partial c}{\partial \kappa_L^*} \right]}^{\substack{4 \\ 5 \quad 6 \quad 7}}} \quad (14)$$

It is clear from (14) 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

$$\frac{\overbrace{[\sigma c(\kappa_L^*)^{\sigma-1} \chi_1 - \sigma + 2] [p - c(\kappa_L^*) + q(\kappa_L^*)]}^{1+3} + \overbrace{2c(\kappa_L^*)^\sigma \chi_2}^2}{\underbrace{w(1 + \kappa_L) [\sigma c(\kappa_L^*)^{(\sigma-1)(1-\sigma)} \chi_3 - \sigma + 1] [p - c(\kappa_L^*) + q(\kappa_L^*)]}_{4+6} - \underbrace{2c(\kappa_L^*)^\sigma \chi_4}_5} > 0 \quad (15)$$

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 2 and 3

In Appendix B we found that

$$\frac{\partial \kappa_L^*}{\partial w_L} = \frac{[\sigma c(\kappa_L^*)^{\sigma-1} \chi_1 - \sigma + 2] [p - c(\kappa_L^*) + q(\kappa_L^*)] + 2c(\kappa_L^*)^\sigma \chi_2}{w(1 + \kappa_L) [\sigma c(\kappa_L^*)^{(\sigma-1)(1-\sigma)} \chi_3 - \sigma + 1] [p - c(\kappa_L^*) + q(\kappa_L^*)] - 2c(\kappa_L^*)^\sigma \chi_4} \quad (16)$$

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., $\frac{\partial \kappa_L^*}{\partial w_L \partial \beta} < 0$.

D Additional Regression Tables

Table D1. Conditional Logit Model, Predicted Distribution: Process Innovation, Baseline

	(1)	(2)	(3)	(4)	(5)
	Process Innovation	Process Innovation	Process Innovation	Process Innovation	Process Innovation
EU8 Share x Post-2004	0.011 ^{***} (3.34)	0.019 ^{**} (2.60)	0.006 [*] (1.71)	0.009 ^{***} (4.09)	0.025 ^{**} (2.58)
EU8 Share	-0.004 (-0.10)	-0.011 (-0.27)	-0.010 (-0.26)		
Skill Share		0.018 (1.70)	0.016 (1.38)	-0.0421 [*] (-2.23)	-0.042 [*] (-2.24)
Wage Change			-0.001 (-0.37)		-0.001 (-0.68)
Output Change			0.000 ^{***} (3.55)		0.000 (0.59)
Observations	8555	8089	8078	8089	8078
Year FE	Yes	Yes	Yes	Yes	Yes
Other FE	No	No	Industry	Firm	Firm

Notes: EU8share is the actual share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a conditional logit model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.

Table D2. Conditional Logit Model, Predicted Distribution: Product Innovation, Baseline

	(1)	(2)	(3)	(4)	(5)
	Product Innovation				
EU8 Share x Post-2004	-0.056 ^{***} (4.47)	-0.088 ^{***} (3.77)	-0.091 ^{**} (2.29)	-0.083 ^{**} (-2.70)	-0.066 ^{**} (-2.69)
EU8 Share	-0.281 [*] (-2.34)	-0.425 [*] (-2.12)	-0.434 [*] (-2.11)		
Log Turnover		-0.590 ^{**} (-3.15)	-0.607 ^{**} (-3.20)	-0.350 (-0.12)	-0.349 (-0.12)
Skill Share		0.466 (1.36)	0.460 (1.35)	0.115 (0.17)	0.159 (0.24)
Wage Change			-0.004 (-0.31)		0.032 (0.63)
Output Change			0.001 (1.25)		-0.004 (-1.24)
Observations	8555	8089	8078	8078	8078
Year FE	Yes	Yes	Yes	Yes	Yes
Other FE	No	No	Industry	Firm	Firm

Notes: EU8share is the actual share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a conditional logit model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.

Table D3. Linear Probability Model, Predicted Distribution: Interactions

	(1)	(2)	(3)	(4)	(5)
	Process Innovation	Process Innovation	Product Innovation	Product Innovation	Product Innovation
EU8 Share x Post-2004	0.016 (0.30)	0.017 (0.27)	0.011 (0.22)	-7.299 (-1.54)	0.008 (0.16)
EU8 Share x Post-2004 x Log Turnover		0.004** (2.71)		0.688 (0.79)	
EU8 Share x Post-2004 x Skill Share	-9.035*** (-4.26)		-4.11* (-2.03)		
EU8 Share x Post-2004 x Process Innovation					0.010* (1.86)
Wage Change	-0.001 (-0.67)	-0.001 (-0.68)	0.001 (0.48)	0.044 (0.79)	0.001 (0.56)
Output Change	0.000 (0.67)	0.000 (0.64)	-0.000 (-1.03)	-0.003 (-1.07)	-0.000 (-1.21)
Observations	8096	8078	8096	8096	8096
Year FE	Yes	Yes	Yes	Yes	Yes
Other FE	Firm	Firm	Firm	Firm	Firm

Notes: EU8share is the actual share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a conditional logit model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.

Table D4. Linear Probability Model, Predicted Distribution: Demand Side Effect

	(1)	(2)	(3)	(4)
	Product Innovation	Product Innovation	Process Innovation	Process Innovation
EU8 Share x Post-2004 x Sold Locally	-1.342 ^{***} (-3.13)		-1.098 ^{***} (-3.78)	
EU8 Share x Post-2004 x Sold Within UK		-1.064 ^{**} (-2.27)		-1.142 ^{***} (-4.24)
EU8 Share x Post-2004	-0.004 (-0.08)	-0.005 (-0.09)	-0.002 (-0.05)	0.009 (0.19)
Wage Change	0.001 (0.43)	0.001 (0.43)	-0.001 (-0.84)	-0.001 (-0.88)
Output Change	-0.000 (-1.46)	-0.000 (-1.43)	0.000 (0.64)	0.000 (0.82)
Observations	8544	8544	8524	8524
Year FE	Yes	Yes	Yes	Yes
Other FE	Firm	Firm	Firm	Firm

Notes: EU8share is the actual share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a linear probability model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Sold Locally is an indicator variable for whether all firm output was sold within 200 miles of production. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.001

Table D5. Conditional Logit Model, 1991 Distribution: Process Innovation, Baseline

	(1)	(2)	(3)	(4)	(5)
	Process Innovation	Process Innovation	Product Innovation	Product Innovation	Process Innovation
EU8 Share x Post-2004	0.097* (2.02)	0.039* (1.87)	0.040 (1.30)	0.032** (2.77)	0.029* (1.96)
EU8 Share	0.654 (0.66)	-5.368*** (-3.41)	-4.398* (-2.35)		
Log Turnover		0.177*** (7.55)	0.173*** (7.35)	0.307** (3.11)	0.310** (3.12)
Skill Share		0.101* (2.30)	0.095* (2.19)	-0.308 (-1.79)	-0.302 (-1.79)
Wage Change			-0.001 (-0.58)		-0.003 (-0.42)
Output Change			0.000*** (3.34)		0.000 (0.85)
Observations	8555	8089	8078	8078	8078
Year FE	Yes	Yes	Yes	Yes	Yes
Other FE	No	No	Industry	Firm	Firm

Notes: EU8share is the actual share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a conditional logit model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.

Table D6. Conditional Logit Model, 1991 Distribution: Product Innovation, Baseline

	(1)	(2)	(3)	(4)	(5)
	Product Innovation	Product Innovation	Product Innovation	Product Innovation	Product Innovation
EU8 Share x Post-2004	-0.981 ^{***} (-8.18)	-0.641 ^{***} (-6.28)	-0.563 ^{***} (-6.51)	-1.07 ^{**} (-3.15)	-1.04 ^{**} (-3.13)
EU8 Share	5.021 (1.93)	0.085 (0.03)	1.300 (0.42)		
Log Turnover		-0.058 (-1.04)	-0.053 (-0.94)	0.539 (1.14)	0.543 (1.14)
Skill Share		0.483 (1.51)	0.523 (1.47)	1.143 (0.71)	1.146 (0.71)
Wage Change			0.007 [*] (2.55)		0.027 (1.43)
Output Change			-0.000 (-0.21)		-0.000 (-0.41)
Observations	8555	8089	8078	8078	8078
Year FE	Yes	Yes	Yes	Yes	Yes
Other FE	No	No	Industry	Firm	Firm

Notes: EU8share is the actual share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a conditional logit model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.

Table D7. Conditional Logit Model, 1991 Distribution: Interactions

	(1)	(2)	(3)	(4)	(5)
	Process Innovation	Process Innovation	Product Innovation	Product Innovation	Product Innovation
EU8 Share x Post-2004	6.456** (3.05)	6.645* (2.41)	6.553* (2.18)	-25.83 (-1.56)	-25.69* (-2.04)
EU8 Share x Post-2004 x Log Turnover			0.006*** (4.25)		1.127 (1.43)
EU8 Share x Post-2004 x Skill Share	-5.55** (-2.93)				
EU8 Share x Post-2004 x Process Innovation					-1.10* (1.86)
Wage Change	-0.003 (-0.44)	-0.003 (-0.43)	-0.003 (-0.44)	0.019 (1.01)	0.052 (1.70)
Output Change	0.000 (0.94)	0.001 (0.91)	0.001 (0.92)	-0.001 (-0.80)	0.001 (0.37)
Observations	8096	8078	8096	8096	8096
Year FE	Yes	Yes	Yes	Yes	Yes
Other FE	Firm	Firm	Firm	Firm	Firm

Notes: EU8share is the actual share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a conditional logit model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.

Table D8. Conditional Logit Model, 1991 Distribution: Demand Side Effects

	(1)	(2)	(3)	(4)
	Product Innovation	Product Innovation	Process Innovation	Process Innovation
EU8 Share x Post-2004 x Sold Locally	-1.16 (-1.20)		-2.40 ^{***} (-4.00)	
EU8 Share x Post-2004 x Sold Within UK		-2.351 [*] (-2.12)		-2.56 ^{***} (-4.88)
EU8 Share x Post-2004	7.810 ^{***} (4.45)	7.938 ^{***} (4.32)	4.175 (1.87)	4.966 [*] (2.31)
Wage Change	-0.000 (-0.06)	-0.000 (-0.05)	-0.006 (-0.66)	-0.005 (-0.62)
Output Change	-0.001 (-1.68)	-0.001 (-1.63)	0.000 (0.85)	0.000 (0.89)
Observations	8544	8544	8524	8524
Year FE	Yes	Yes	Yes	Yes
Other FE	Firm	Firm	Firm	Firm

Notes: EU8share is the actual share of individuals working in a TTWA who were born in EU8 countries in the 1991 Census. We estimate a linear probability model with a difference-in-differences approach. The controls for wage changes and output changes are changes in these variables over the pre-period, between 1997 and 2003. The most restrictive specifications include firm and year fixed effects. Sold Locally is an indicator variable for whether all firm output was sold within 200 miles of production. Standard errors are clustered at the TTWA level. t statistics in parentheses. * p<0.05, ** p<0.01, *** p<0.001