

PROCESSING IMMIGRATION SHOCKS: FIRM RESPONSES ON THE INNOVATION MARGIN*

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

March 30, 2020

Abstract

The extent to which firms respond to labor supply shocks has important implications for local and national economies. We exploit firm-level panel data on product and process innovation activities in the United Kingdom and find that the large, unanticipated, low-skill labor supply (immigration) shock generated by the 2004 expansion of the European Union to Eastern European countries increased process innovation and reduced product innovation, with overall innovation activity going up. 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 is consistent with a demand side effect generated by the labor supply shock.

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

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

*We thank many seminar participants along with Laura Giuliano and Ethan Lewis for helpful comments. Two anonymous referees made suggestions that greatly improved this work. Montresor acknowledges the support of an Economic and Social Research Council PhD Studentship. We thank the University of Essex for facilitating continual access to the secure data and providing Greg Wright with a visitor travel grant.

[†]Gray: Department of Economics, University of California Merced. Email: rgray6@ucmerced.edu.

[‡]Montresor: Department of Economics, University of Mannheim. Email: montresor@uni-mannheim.de.

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

1 Introduction

The impact of immigration on local economies is the subject of a large literature. Among other important findings, there is strong evidence that even large inflows of immigrant workers produce little impact on native employment rates and wages (see Card [13] for a discussion). Several mechanisms have been suggested to explain this. For instance, it has been shown that native and immigrant workers in many cases have complementary skills, even within low-education skill categories, which leads to productivity gains when these workers are used together (Peri and Sparber [43]). 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 [20] and Lewis [37]) while, possibly at the same time, adjusting their capital stock or adopting new technologies in response to the labor supply shock (e.g., Lewis [36], Lewis [37], Lafortune et al. [35] and Ottaviano and Peri [41]), both of which may mitigate any local wage and employment effects. More generally, Acemoglu [1] argues that firms will respond to changing skill supplies and premia by re-optimizing over the technologies used. In this paper, we explore two channels of firm response to labor supply shocks, namely, changes arising from process innovations and changes in firms' products due to product innovations, both of which may impact the distribution of output and employment within and across firms, with subsequent welfare consequences for workers.

There is little extant literature that separately relates labor supply shocks to process or product innovations, though Lewis [37] outlines a range of potential firm responses, some of which are consistent with the mechanisms we present here.¹ This gap in the literature is somewhat surprising, as process innovations have been shown to be a key aspect of the firm response to technology adoption² and international trade,³ suggesting that labor supply shocks may also induce these types of responses. Here we provide evidence to this effect.⁴ In addition, we provide new evidence on the role that labor supply shocks play in firms' decisions regarding optimal product scope and quality (product innovations).

Generally speaking, the welfare consequences of process and product innovations are likely to be of first order importance, via well-known channels. Falling production costs due to process innovations will typically lead to price reductions and corresponding welfare gains. Relatedly, Cortes [14] shows that immigrants reduce prices of immigrant-intensive output, though in that case it is by reducing wages, and hence production costs. The channel we explore in this paper

¹A paper that is similar to ours is Maré et al. [38] who study the relationship between innovation and immigration in New Zealand, finding that the relationship between immigration inflows and innovation outcomes is a function of firm characteristics. We also find that firm responses are mediated by firm characteristics, as we discuss further in our results section. Another relevant paper is Mazzolari and Neumark [40] who find that low-skill immigration leads to a more homogenous retail sector.

²See Markus and Robey [39] for an early discussion and Bloom et al. [10], Bloom et al. [11] and Gaggl and Wright [27] more recently.

³See, e.g., Antras et al. [4] and Antràs and Rossi-Hansberg [5].

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

is an additional potential mechanism through which production costs and prices may fall due to increased immigration, in this case via increased process innovation. At the same time, increased process innovation may straightforwardly raise worker incomes by increasing firm productivity (e.g., see Huergo and Jaumandreu [32] or Bloom et al. [9]). And finally, product innovations should lead to a fall in the price index, and will thus increase welfare, by expanding the range of available product varieties or reducing quality-adjusted prices (e.g., see Feenstra [25] or Eizenberg [22]).

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

We exploit the expansion of the European Union (EU) to Eastern Europe in 2004 as a differential, and large, shock to the supply of low-skill labor across UK local labor markets. This large inflow of immigrants to the UK was mostly unanticipated since, historically, the UK was a low-immigration country and expert predictions of the potential inflows due to the expansion were quite small. Using firm-level panel data on product and process innovation activities, we estimate difference-in-differences specifications that produce consistent results. We first find that the immigration shock increased process innovation, on average. Noting that firms are likely to respond to supply shocks very differently, we then explore heterogeneity in the response, finding that the response was increasing in firms’ low-skill production intensity as well as firm size. We also find that product innovation *fell* in response to the migration, an effect that is greater within low-skill intensive firms but smaller for firms whose output is sold locally. We interpret this last finding as evidence on the importance of a demand side effect.⁵

Existing research on the impact of labor supply shocks on innovation has typically focused on the impact of high-skill immigrants on patenting and overall knowledge creation. For instance, Stuen et al. [44] exploit a shock to the supply of foreign doctoral students in science to measure their impact on knowledge creation in the U.S., finding a large, positive impact. Hunt [33] 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, but do not crowd out native innovation. Kerr [34] provides a comprehensive review of studies looking at skilled immigration and innovation outcomes as proxied by patenting and firm starts. For the U.S., 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-

⁵Bodvarsson et al. [12] find support for a similar demand side effect in the context of the Mariel Boatlift.

born 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. [42]).

The paper is organized as follows. Section 2 presents some stylized facts. In Section 3 we present a conceptual framework to guide the empirical analysis. Section 4 describes the data. Section 5 introduces the empirical specifications and identification strategy. Section 6 discusses the results and Section 7 concludes.

2 Stylized Facts

2.1 EU8 Immigration to the UK as an Exogenous Shock

This paper exploits a large, policy-induced shock to the relative supply of low-skill labor across UK travel to work areas (TTWAs) in the form of the expansion of the EU in 2004. TTWAs are standardized UK local labor markets, a geographic unit developed by the Office of National Statistics (ONS). In short, they are defined in order to cover both metropolitan areas as well as their commuter suburbs.⁶ The expansion brought in eight Central and Eastern European countries: Czech Republic, Estonia, Hungary, Latvia, Lithuania, Poland, Slovakia and Slovenia, with the majority (75%) of UK immigrants from these countries over the 2004 to 2006 initial wave coming from Poland. We refer to these countries collectively as the EU8 countries. Though EU8 citizens immediately enjoyed free movement across EU countries, their access to most labor markets was restricted during a seven-year phase-in period. The exceptions were Ireland, Sweden and the UK who granted immediate access, the result of which was a large inflow of migrants into these countries. The UK restricted their access to benefits, with the result that migrants can be expected to be fully engaged in the labor market during this period. Blanchflower and Lawton [8] bear this out, showing that EU8 immigrants from 2004 to 2008 were 13 percentage points more likely to be working compared to natives and 5 percentage points more likely than pre-2004 immigrants from EU8 countries. Figure 1 depicts the long-run trend in immigration to the UK, indicating that 2004 was a significant turning point. In Figure 2 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 [23]).

Most important for the purposes of our research design is that the magnitude of the inflow to the UK was largely unanticipated. Negotiations over the terms under which the countries would enter the EU concluded 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 5,000-13,000.⁷ At the time that the document was published it was not known with certainty whether or not Germany would

⁶Formally, 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”.

⁷See Dustmann et al. [18].

impose labor controls on the new accession countries, and the authors emphasized that if Germany maintained labor controls then some of the expected flow of up to 200,000 migrants might divert to the UK. From the perspective of UK firms, the low anticipated flows into the UK were likely believable 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 that the UK decision to maintain open borders was taken solely by the UK government, with limited consultation from labor market actors, in contrast to the debates that occurred in other large European countries. Those countries did not fully open their borders for labor migration until 2011, after the end of our sample period.

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.⁸ According to Dustmann et al. [19] the average hourly wage over the period 2004-2009 for men from EU8 countries was £6.81 while it was £11.91 for native-born men. Blanchflower and Lawton [8] show that the most common occupations for EU8 workers up to 2008 were process operative and warehouse operative. This suggests that the EU8 expansion significantly changed the labor force composition in areas that received large numbers of these migrants, skewing it more towards low-skill labor. To the extent that there was downward pressure on the average low-skill wage, this will have generated a productivity gain for firms who employed these workers, and relatively more so for firms who used low-skill labor relatively intensively.

3 Conceptual Framework

The primary focus of the paper is on the extent to which a large, low-skill labor supply shock induces product and process innovation investments within UK firms. Here we present a description of the possible channels of local response of firms to a labor supply shock in their market, with a specific emphasis on the role played by product and process innovation responses.

We assume there are many distinct UK labor markets, each containing $i = 1, \dots, I$ firms that produce output Y^i . Production requires the input of several factors, including low-skill labor. Denoting L as the supply of low-skill labor in some labor market, full employment in that market requires that

$$L = \sum_{i=1}^I c_L^i(A^i, W)Y^i, \quad (1)$$

where $c_L^i(A^i, W)$ is the low-skill unit labor requirement associated with firm i , A^i is a vector of factor-augmenting firm production technologies, and W is a vector of local factor prices. We can totally differentiate (1) to get

⁸This was despite their higher average education level, suggesting that either their origin country education was of low quality or was simply not highly relevant for jobs in the UK (see Dustmann et al. [19]).

$$dL = \sum_{i=1}^I Y^i c_{LA^i}^i dA^i + \sum_{i=1}^I c_L^i dY^i + \sum_{i=1}^I Y^i c_{LW}^i dW, \quad (2)$$

where $c_{LA^i}^i$ is a vector of derivatives of firm i 's unit factor requirement with respect to firm production technologies and c_{LW}^i is a vector of cross-price derivatives of the unit factor requirement with respect to local factor prices. In short, the condition states that changes in the local supply of low-skill labor (dL) must be offset by changes in individual firms' demand for low-skill labor (unit low-skill labor requirements) resulting from some combination of a change in firm technology, a shift in firm output, or a change in local factor prices.

We take each of the three terms on the right hand side of (2) in turn. With respect to the first term, we note that Acemoglu [2] considers the case of a general production technology and shows that under mild assumptions an increase in the supply of a factor always induces technological change that is biased toward that factor, thereby raising demand for the factor (what he calls *weak absolute equilibrium bias*). In our case, we assume that process innovations are among the set of technological changes (dA^i) available to the firm. As a result, a low-skill labor supply shock may induce firm-level process innovations that will help to equilibrate the local labor market by raising the demand for low-skill labor within the firm.

With respect to the second term in (2), we can assume that firm output, Y^i , is comprised of the output of a set of N products, indexed by j , whose average output is \bar{y}^i , so that $Y^i = N\bar{y}^i$. This implies that induced changes in firm output may arise due to changes in the average output per product or alternatively via a change in the number of products. More specifically, the Rybczynski theorem states that the composition of output will shift toward firms that are relatively intensive in the now-more-abundant factor, at the expense of other firms' output.⁹ The net result will be to raise aggregate output as well as the total demand for that factor (e.g., see Hanson and Slaughter [29]). Again, in our case the net rise in output induced by the low-skill labor supply shock will come from either 1) an increase in the average output of products produced by low-skill-intensive firms, or else 2) an increase in the number of products produced by low-skill-intensive firms (i.e., product innovation).

The third term in (2) represents the change in the price of low-skill labor relative to the price of other factors. This has both direct and indirect effects. The direct effect is to alter the demand for low-skill labor (by a magnitude of $Y^i c_{LW}^i dW$, i.e., the third term in (2)). The indirect effect is to alter the opportunity cost of engaging in activities that are intensive in other factors. Thus, to the extent that process or product innovation are high-skill-intensive activities the downward pressure on the low-skill wage due to a positive low-skill labor supply shock will *reduce* firm incentives to innovate, thereby counteracting the effects in favor of process and product innovation described above. As a descriptive exercise we note that the cross-firm correlation between a product innova-

⁹In a model in which products are also heterogeneous in their unit skill requirements the composition of output *within firms* will of course also shift toward production of products that are relatively intensive in the now-more-abundant factor.

tion indicator and the high-skill employment share (defined in Section 4 below) is 0.58, while the correlation between a process innovation indicator and the high-skill share is 0.19. In other words, both innovation types are positively correlated with skill, though product innovation much more so. This suggests that the downward pressure on the low-skill wage induced by the labor supply shock will reduce firms' incentives to innovate as the opportunity cost of innovation rises, though whether this is enough to offset the rise in innovation directly induced by the first two channels in (2) is an empirical question.

We note that this analysis has focused on the supply side response to a labor supply shock. However, various models have been proposed in which demand side factors can play a key role in firms' decisions about whether to alter their product scope. In these models a rise in the local labor supply also generates an expansion in the size of the market for firms' output. Whether this leads firms to increase or decrease their product scope (i.e., increase or decrease product innovation) ultimately depends on the details of the demand system. For instance, Feenstra and Ma [24], Eckel and Neary [21], and Dhingra [15] describe models in which "cannibalization" dominates, such that the increased competition induced by the rise in market size leads firms to *reduce* their product scope. In contrast, Di Giovanni et al. [16] show that an immigration-induced rise in market size can lead firms to *increase* their product scope.¹⁰ We provide some new evidence on the relative strength of these channels in our empirics.

We also note that this analysis assumes perfectly segmented labor markets, whereas UK labor markets are likely to be at least partially integrated. When labor markets are not perfectly segmented then we must also allow for national labor market clearing so that out-migration of workers from some labor markets are offset by in-migrations elsewhere. Conceptually this implies that the general equilibrium change in labor supply in any market is endogenous to the right hand side of (2), as workers reallocate across markets in response to subsequent changes in firm technology, output and local factor prices.

Finally, process and product innovation will be jointly determined in equilibrium, such that intended or realized process innovations may incentivize firms to alter their level of product innovation. On the one hand, process innovations may increase firm productivity, and therefore increase the profitability of new products (increasing product innovation). On the other hand, there may be a within-firm tradeoff between process and product innovation activity in the short run, perhaps due to the scarcity of shared inputs (e.g., qualified engineers). In this case a rise in process innovation may reduce the likelihood of product innovations. We consider these possibilities in the empirics.

Taken together, the channels described above imply that the effects of a low-skill labor supply shock on product and process innovation are ultimately empirical questions. We interpret our findings with these channels in mind.

¹⁰Interestingly, in an international trade context Hottman et al. [31] find that cannibalization is an important channel of adjustment to trade shocks. In contrast, Di Giovanni et al. [16] find evidence that the market size expansion due to immigration inflows led firms to increase their product scope.

4 Data

In Section 5 below we explore the innovation response of UK firms to a labor supply shock in their local labor market, which we define as a UK Travel to Work Area (TTWA).¹¹ To do this, we exploit data on employment from the UK Quarterly Labour Force Survey (QLFS) across 243 UK TTWAs.¹² The QLFS is a quarterly sample of workers that includes a variety of work-related and demographic information, including the worker’s country of birth. Our identification strategies also exploit cross-sectional variation in the EU8 immigrant share of the population and economic outcomes from the 1991 Census. Summary statistics for our primary variables of interest are reported in Table 2.

Our dependent variables exploit firm-level panel data on innovation activities from three waves of the Community Innovation Survey (CIS), covering the period 2002 to 2008.¹³ 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 changes in production structure during the previous three years. It is conducted every four years, such that we exploit survey responses regarding firms’ innovation activities at three snapshots: 1) between 2002 and 2004 – the period prior to the EU8 accession – as well as 2) between 2004 and 2006 and 3) 2006 and 2008. The nature of the timing of the survey requires two comments. First, there is an overlapping year in each survey, however this is inconsequential given the binary nature of our main outcome variables.¹⁴ 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 did not engage in product innovation in 2004 (and, of course, did during 2002-2003).

Second, we must address the fact that 2004 is assigned to the pre-period, even though one of our identification strategies exploits variation due to the EU enlargement that occurred on May 1st 2004. As a result, any response by firms from May through December of 2004 due to the immediate inflow of immigrants from EU8 countries will be allocated to our pre-period control group. We note that this will work against finding an effect due to the EU8 accession – i.e., it will bias our results toward zero. Figure 2 documents the trend in EU8 inflows beginning in 2004. We can see that there was indeed an immediate uptick in EU8 migration to the UK beginning in June, 2004, however the vast majority of the inflow occurred after December 2004. Furthermore, given that the government vastly underestimated the projected magnitude of the immigrant inflows, suggesting there was little anticipation of the magnitude, it is likely that any firm response to the labor supply shock lagged the EU expansion.

Finally, the CIS consists of a stratified sample of approximately 28,000 firms with more than 10 employees. For the period we are interested in, 2002-2008, the CIS has a panel dimension consisting

¹¹We use the 2001 ONS definition of a TTWA.

¹²Our replication package can be found on Mendeley at Gray and Wright [28].

¹³Note that we use the population-representative panel dimension of the data in part because our specifications use firm fixed effects. The unbalanced dataset would drop firms observed in only a single year and the resulting sample would be unrepresentative.

¹⁴We 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.

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.¹⁵ It further asked for firms' spending on R&D, and the objectives of these innovation activities. These questions regarding whether firms actually did innovation may be a more direct measure than the traditional patent data used in the literature, which measure invention rather than innovation. Similarly, we would not want to rely on R&D expenditure entirely because not all expenditures will successfully lead to implementation of new products or processes.

Over the three waves of the CIS 18 percent of the firms report product innovation in each period and 23 percent report process innovation in each period. Across the first and second waves, 38 and 44 percent of firms switched their reported engagement in product and process innovation, respectively, and 39 and 49 percent across the second and third waves. The fact that there are firms who change their innovation decisions during the period is crucial to our analysis and we will estimate how much of their change seems to be driven by the labor supply shock. More generally, the intensity of the treatment (2004 EU-8 share) across TTWAs was uncorrelated or weakly correlated with the initial share of skill or with pre-period employment growth (see Appendix Figure A1). In addition, EU-8 migrants were disproportionately located in London and, to a lesser extent, the South East of the UK (see Table 5). In a robustness check we remove London from the sample and reproduce our main specifications.

In Figure 3 we plot the unconditional change in the share of EU8 immigrants in a TTWA cell over the period 2004-2008 against the mean change in process (Panel A) and product (Panel B) innovation across firms over the period (see Section 4 for a description of the immigration and innovation datasets). The plots are clearly only suggestive, but they indicate a positive correlation between the group of relatively low-wage EU8 migrants and the extent of process innovation and a negative correlation between these migrants and the extent of product innovation. This suggests that process innovation may have been induced by the increased availability of low-skill labor, which may have led to a reorganization of production consistent with the weak absolute equilibrium bias outlined by Acemoglu [2] (see Section 3). On the other hand, product innovation is seemingly reduced by this inflow of workers. In the context of the conceptual framework, this is consistent with a dominant role for opportunity cost relative to output shifts due to the supply shock. In other words, the increased relative abundance of low-skill labor would have increased the opportunity cost of product innovation to the extent that it is a high-skill intensive activity, and this may have offset the rise in low-skill intensive output (and thus product innovation) generated by the low-skill labor supply shock. We address these issues more carefully in the empirical section below.

¹⁵This is paraphrased from the 2008 CIS documentation.

What is Process Innovation?

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

As we can see from the table, the strongest (conditional) correlates with process innovation are “Improvements in Production Flexibility”, “Improvements in Production Capacity” and “Reduce Per Unit Costs”. Additionally, there is virtually no correlation with contemporaneous capital investment or capital intensity. Across both specifications “Reduce Per Unit Costs” is highly significant, which is consistent with a standard theoretical treatment of process innovation as a simple scaling of unit costs (see, e.g., Dhingra [15] or Duranton and Puga [17]).

5 Specifications & Identification

In this section we outline our identification strategy, describe the underlying source of identifying variation, and perform a range of tests of the identifying assumptions.

5.1 Baseline

In our baseline specification we exploit the discontinuous inflow of immigrants arising from the 2004 EU8 expansion, described in Section 2 above. To exploit the discontinuity, we begin by considering the following difference-in-differences specification:

$$INN_{iat} = c + \beta_1 [POST_t \times EU8SHR_{a,2004}] + \alpha_t + \gamma_i + \epsilon_{iat} \quad (3)$$

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

from projecting the untreated potential outcome (i.e., when the treatment is zero) on the control variables.

The differential extent of the treatment is defined by the cross-sectional variation in the share of EU8 immigrants in a TTWA at the beginning of the period, 2004. This approach exploits a version of the “ethnic enclave” design commonly associated with Altonji and Card [3]. 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 UK. The share of Hungarians in an area in 2004 should therefore serve as a useful predictor of Hungarian settlement between 2004 and 2008.

Figure 4 suggests that this is the case. The x-axis reports the share of immigrants in a TTWA from each of the EU8 origin countries in 2004, while the y-axis reports the percentage point rise in the share of immigrants from each country in a TTWA between 2004 and 2008. In other words, each observation represents a single TTWA-by-EU8-origin-country cell (243 TTWAs by 8 origin countries), and we relate the 2004 level to its subsequent growth. We see that areas with a higher share of immigrants from a particular EU8 country in 2004 on average saw a larger rise in immigration from that country relative to other immigrant groups in subsequent years.¹⁶

We estimate each version of (3) with both a linear probability model (LPM), our main specification, as well as a conditional logit (CL) specification, reported in Appendix Tables A1-A5. Since in these specifications the treatment varies across TTWAs in the cross-section, we cluster standard errors at the TTWA level throughout. All regressions are weighted by TTWA employment though the qualitative findings are not sensitive to these weights, including the use of firm-size weights. In Section 6 below we discuss these estimates.

5.2 Using the Lagged EU8 Distribution

The exclusion restriction with respect to specification (3) may be violated, since innovation outcomes within a local labor market over the 2004-2008 period may be driven by shocks that also drove immigrants to that area in the years leading up to 2004.

We therefore estimate specification (3) both using the 2004 distribution (as in (3)) and also using a lagged EU8 share variable, $EU8SHR_{a,1991}$ in place of the 2004 distribution, reflecting the share of EU8 immigrants in a TTWA in 1991.¹⁷ 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

¹⁶We note that the slope is statistically significant at the one percent level (p-value 0.003), with a partial F-Stat of 12.

¹⁷Bell et al. [7] use the 2001 Census distribution of immigrants from EU8 countries to implement a Bartik-type procedure in an attempt to isolate the exogenous component of immigration across locations within the UK. Foged and Peri [26] also apply a difference-in-differences strategy that is similar to ours. Specifically, their treatment is an interaction between an inflection point in which immigration rapidly increased and an indicator for whether a region was in the highest or lowest quartile of predicted inflows based on a historical refugee dispersal policy.

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 discuss and test in the next section.

Given these specifications, the two key issues with respect to identification are: what are the sources of variation driving the 1991 and 2004 cross-section EU-8 immigration shares? And, is this variation unrelated to changes in firm outcomes over the 2004-2008 period (except through immigration inflows)?

Migrant Network Formation

The first question is particularly important given that the countries that joined the EU in 2004 had been behind the Iron Curtain for much of the late twentieth century, so that there was not a continuous flow of migrants from these countries over that period. Nevertheless, strong connections were maintained between pre-Cold War migrants to the UK and their homelands, in part as a response to the adversity that was faced by family members who remained in Eastern Europe. For instance, the Poles – by far the largest migrant group to the UK following the 2004 EU Accession (see Table 4) – formed large communities around the UK as of 1951 (see Holmes [30]), in large part due to a large-scale resettlement plan for Poles following World War II.¹⁸ According to the 1951 Census there were 152,000 Polish-born UK residents, constituting the second largest UK immigrant group. These immigrants maintained close connections with family and friends during the Cold War and beginning in 1956 some travel was allowed between Poland and the UK, which facilitated the maintenance of networks. Thus, for members of the Eastern European diaspora the desire to remain connected to their homelands was often heightened during the Cold War.

The fall of the Iron Curtain, between 1989 and 1992, led to migration that was largely driven by these networks, and is the period in which we set our baseline cross-section of EU8 immigrants.¹⁹ Future immigration, for instance due to the 2004 EU Accession, was then also likely to be in part driven by these established settlement patterns. And the 1991 immigrant distribution is indeed predictive of the 2004 distribution. We confirm this formally, finding a strong and significant correlation between the 1991 and 2004 immigrant distributions, with an F-Stat of 35, reflecting a persistence in the geography of EU8 migrants over two decades. In Table 4, Row 3 we report the correlation for each of the EU8 immigrant groups, where we see that for most groups the correlation is strong. In fact, the correlation is strongest for Polish immigrants, who subsequently experienced the largest growth in population within the UK following the 2004 Accession. Ultimately, EU8 migrants settled throughout the country post-2004, as Table 5 depicts.²⁰

¹⁸Known as the Polish Resettlement Act of 1947 it offered British citizenship to Polish troops living in Britain after the war. Most settled in London, Swindon and the industrial areas of the North.

¹⁹Note that it would be ideal to use the 1951 distribution of EU8 migrants across UK labor markets – which was almost entirely the product of post-war resettlement – as our baseline cross-section, but these data are not available.

²⁰These data come from the UK Worker Registration Scheme which was explicitly set up to monitor the inflows of EU8 migrants into the UK.

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

Table 5: EU8 Migrant Population by UK Region and Year

We note here that since London is clearly the primary hub for immigration to the UK it is important that we also run our baseline empirical analysis with London removed. Appendix Tables A6 and A7 present the results from this exercise (finding very similar results).

Testing the Common Trends Assumption

Networks were therefore very likely to be drivers of the persistence in the immigrant distribution between 1991 and 2004. However, this persistence may also reflect a correlation between economic shocks leading up to 1991 that drove immigration to specific UK locations (or within-UK migration of existing EU8 residents) and that were also correlated with innovation outcomes in the later period. For instance, regions receiving positive economic shocks during the late nineteen eighties may have attracted EU8 immigrants (or internal EU8 migrants), and may have remained desirable locations to live over the subsequent two decades, for reasons that also led to greater firm innovation post-2004. To address this, we first ask whether the immigrant distribution in 1991 is correlated with various firm outcomes across TTWAs over the period 2004-2008. These correlations help to address (though they do not definitively establish or disprove) the common trends assumption by asking whether there are common correlates of our dependent and independent variables. In other words, we ask whether there are underlying economic trends driving the distribution of EU8 migrants in 1991, and whether these trends are sufficiently persistent to drive firm outcomes over 2004-2008. It is important to note that the existence of correlations would not necessarily indicate that our estimates are biased (e.g., the distribution of immigrants across TTWAs may be dominated by urban areas, but still may be exogenously allocated within that set of areas), but it will help to identify variation that we would be wise to control for in our specifications.

Table 6 reports the results of separate OLS regressions in which the dependent variable is growth over the 2004-2008 period in: employment, output per worker, exports, imports, or capital investment. We also include a 2004 urban-rural indicator and TTWA distance to London as dependent variables. The regressors are, separately, the 1991 immigrant distribution and the 2004 immigrant distribution, which we provide simply for additional context. We see that the takeaway is that both the 2004 and 1991 immigrant distributions are skewed toward urban areas that are globally integrated. Though this does not necessarily indicate a violation of the common trends assumption in our variables of interest, in our baseline regressions we will include a specification that controls for an urban indicator as well as log exports and imports, each interacted with the post-2004 treatment indicator. We note that the controls for differential trade exposure across periods may be particularly important due to the fact that the EU accession may have reduced trade costs between the UK and Eastern Europe, alongside the reduction in migration costs. Column 1 of Table 6 also shows the relationship between the EU8 immigrant shares in each of our baseline years, 1991 and 2004, and the change in EU8 share over the treatment period, 2004-2008. These coefficients are similar to a “first stage”, were we running an instrumental variables regression instead of the reduced form. They show similar magnitudes for each of the years and the implied instrumental variables estimates are similar to our actual estimates reported below.

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

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

$$INN_{iat} = c + \lambda_1 [POST_t \times EU8SHR_{a,2004} \times X_{ia}^{init}] + Z_{iat} + \alpha_t + \gamma_i + \epsilon_{iat} \quad (4)$$

where Z_{iat} is a vector of the required two-way interaction terms and X_{ia}^{init} is the value of the interaction term of interest in the initial period, 2004. The use of initial-period values should mitigate the potential endogeneity of these measures.²¹ We then estimate the reduced form (4) again but replace $EU8SHR_{a,2004}$ with $EU8SHR_{a,1991}$.²²

First, we estimate specifications in which product innovation is the dependent variable and in which the treatment is interacted with the (pre-period) *process* innovation indicator. As noted in Section 3, equilibrium requires that both process and product innovation are jointly determined. As a result, the direct (cost-reducing) productivity gains associated with process innovation may raise optimal product innovation by raising the profitability of new products. This specification provides evidence on this channel.

²¹Note that the individual terms from the interaction are absorbed in the firm fixed effects.

²²Recall that each of these values is multiplied by 100 for ease of interpreting the estimates, as in our baseline specification.

We next explore the role of firm size and firm heterogeneity in worker skill in the magnitude of the process innovation response. Throughout, we proxy the low-skill share with the share of employees without a college degree in science or engineering subjects in the pre-period ($LowSkillShare_{ia}^{pre}$).²³ We then proxy firm size with firm revenue in the pre-period ($Revenue_{ia}^{pre}$).

In our final specification we directly explore the demand side impact of the labor supply shock. We noted in Section 3 that existing theory and empirical evidence is mixed as to the predicted response of product innovation due to demand side channels. Here we exploit information on the location of firm sales, namely a set of indicators for whether the firm sold all of their output 1) locally (within 200 miles), 2) within the UK only, or 3) both in the UK and internationally. We estimate a specification with two treatment terms: the first term interacts the treatment intensity term (Post-2004 x EU8 immigration share) with an indicator for whether the firm initially sold all of their output locally ($LocalSales_{ia}^{init}$) and a second term that interacts the treatment with an indicator for whether the firm sold all of their output within the UK ($UKSales_{ia}^{init}$). We again do this using the 2004 and, separately, the 1991 EU8 distributions. 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 product innovation effects due to this channel to be concentrated locally as well, and to a much lesser extent UK-wide (relative to UK firms that sell both domestically and abroad, i.e., the omitted group). We note that a key identifying assumption is that any supply side effects should operate independently of whether firm sales are local versus UK-wide, such that the estimated effects arise from the demand side only.

6 Results

Below we discuss the results of a set of LPM regressions. The results are robust to estimation via conditional logit, which are presented in Appendix Tables A1-A5.

6.1 Product Innovation Estimates

The conceptual framework in Section 3 indicates that the effect of the EU8 immigrant labor supply shock on product innovation is ambiguous and depends on the relative strength of the output growth associated with the reallocation of output due to the shock and the extent of substitution away from product innovation due to its high-skill intensity. Additionally, the increased local market size due to the shock may be a force to increase average firm product scope, or may lead to product “cannibalization”, depending on the nature of demand. With these theoretical ambiguities in mind, Table 7 presents OLS results based on (3) where the dependent variable is a binary indicator for

²³Since the relevant “skill” that we are interested in is the skill required to develop and implement new product or process innovations, this measure of science and engineering education is likely an accurate measure of this type of human capital. We also note that this measure of skill is correlated with firm productivity. More specifically, an additional percentage point in the high-skill share within a firm is associated with an increase in output per worker of £4054, on average. With this in mind, we attempt to control for differential changes in productivity over the treatment period in our regression specifications.

whether the firm engaged in product innovation in either period, and we note again that the pre-treatment period covers 2002-2004. Columns (1)-(3) exploit the 2004 distribution of EU8 immigrant shares while columns (4)-(6) exploit the 1991 distribution. Throughout our empirical analysis we progressively increase the strictness of the specifications, in our final two specifications adding firm fixed effects. In our strictest specifications (columns 3 and 6) we also interact the treatment term with several variables in order to control for other channels that may influence firm innovation arising from the EU accession (or other contemporaneous events), in line with our discussion of threats to common trends in Section 5.2. Time fixed effects are included throughout.

We find that the effect of the 2004 accession on product innovation is negative and statistically significant across all specifications. The magnitude of the coefficient ranges in size from -0.057 to -0.521, with the smallest estimates associated with specifications using the 1991 EU8 distribution which identifies the historical immigrant enclave-driven effect. Taking the specification in column (6) using the 1991 distribution and with firm fixed effects and controls as preferred, and applying the coefficient on the treatment variable of -0.057 – the estimate indicates that a one percentage point increase in the EU8 immigration share led to a 5.7 percentage point drop in the product innovation rate from the 2004 level. The average observed rise in the EU8 immigrant share across TTWAs due to the EU8 expansion was just over half of one percentage point (see Table 2), which would therefore be associated with a 3.42 percentage point fall in the product innovation rate. On aggregate, product innovation rose three percentage points between 2004 and 2008 (see Table 2), indicating a significant, and countervailing, role for the 2004 immigration shock.

We can also place an upper bound on the economic magnitude by combining these estimates with the innovation premia reported in Table 1. Again, the estimates in Table 1 are simply OLS correlations between the innovation indicators and firm outcomes (with firm fixed effects in some specifications), and are likely biased upwards. We therefore take them as likely upper bound estimates. If we combine the estimates in Table 7, column (6) that relate EU8 immigration to product innovation with the Table 1 correlations, we find that the effect of the observed rise in the EU8 immigration share (of just over half of a percentage point) is associated with a $3.42 \times 1.19 = 4.07$ percent reduction in firm revenue, a $3.42 \times 1.001 = 3.42$ percent reduction in employment and a $3.42 \times 0.672 = 2.30$ percent reduction in R&D investment. Again, these effects are upper bounds and, in addition, may be offset by increased process innovation, which we explore in the next section.

In Table 10 we explore heterogeneity in the product innovation response based on specification (4). The relevant coefficients are those on the triple interaction terms in which we ask if the product innovation response is increasing or decreasing in the share of low-skill workers in the firm, or increasing or decreasing when the firm does, or does not, engage in process innovation. In columns (1) and (4) we see that the negative product innovation response is greater for firms that are intensive in low-skill workers – i.e., firms that are relatively intensive in the now-more-abundant factor are even less likely to engage in product innovation relative to other firms following the labor supply shock. In terms of the economic magnitude, moving from the mean level of the low-skill

share (86 percent) to all low-skill workers (100 percent) would have doubled the (negative) product innovation response, indicating an important role for firm factor intensity. Finally, in Columns (2) and (5) we see that the negative product innovation response is mitigated in large firms, implying that large firms are more insulated from local supply shocks.²⁴ For instance, since large firms both serve more markets and are less dependent on any specific market for inputs, relative to small firms, they may be less responsive to any particular labor market supply shock.

Next, in Columns (3) and (6) we interact the treatment with an indicator for whether the firm engaged in process innovation during the pre-period.²⁵ The results suggest that the negative product innovation response was made more negative for firms who were engaged in process innovation, suggesting a margin of substitution between these two types of innovation in the face of the shock.²⁶ This is interesting in light of the fact that there is an overall positive correlation across firms in the probability of engaging in product and process innovation (if you do one, you are more likely to do the other). One possibility is that the firm faces a tradeoff between the two innovation types in the short run, but not the long run, which could be driven by shared, and scarce, inputs into each innovation activity. To the extent that the firm is constrained in its use of a particular input in the short run – for instance, skilled engineers – this would explain substitution across innovation activities in the face of a shock, while still rationalizing a positive correlation between the activities in the long run (as reflected in the overall correlation between the activities in the cross-section). We note that these results are somewhat weak and so should be interpreted with caution – the estimates are not significant when using the 1991 EU8 distribution, and only significant at the 10 percent level when using the 2004 distribution.

Table 11 then asks 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), relative to the omitted group (who sell both within the UK and internationally). This specification explores the interaction between EU8 immigrants and local demand for firm output in driving the product innovation decision. The estimates using either distribution of EU8 immigrants (2004 or 1991) indicate that the negative product innovation response is in fact mitigated when the firm only sells locally, consistent with a mitigating effect coming from the demand side.²⁷ In the context of the conceptual framework (Section 3) this provides evidence that the low-skill labor

²⁴We should also point out that the “baseline” interaction between the EU8 share and the Post-2004 dummy in row 4 of Table 10 is positive, the opposite of the main effect reported in Table 7. This is due to non-linearity in the effect of revenue on the product innovation response and seems to be driven by the lowest-revenue firms responding *positively* to the shock in the product innovation dimension.

²⁵We note that we restrict this measure to the pre-period in order to isolate variation that is likely to be exogenous to the treatment, so the results should be interpreted as reflecting the impact of product innovation on initially more or less process-intensive firms. We also note that the reported, separate effects of EU8 immigration on product and process innovation imply that they are substitutes, since the effects are of opposite sign. So the interaction exercise here simply confirms this.

²⁶We also performed this exercise for the reverse scenario: looking at the effect of pre-period product innovation on subsequent process innovation, where we also find substitution between the two. Specifically, the positive effect of the labor market shock on process innovation is lessened for higher initial levels of product innovation.

²⁷We note that we also perform this exercise in the context of process innovation, but there we find no statistically different-sized effects across firms that sell locally versus non-locally.

supply shock is a force for expanding product scope via demand-side channels, in line with evidence from, e.g., Di Giovanni et al. [16].

Finally, in Appendix Table A6 we present the baseline product innovation results (similar to Table 7) with London removed from the analysis, and we find that the coefficients are somewhat smaller, but otherwise consistent with the baseline results.

6.2 Process Innovation Estimates

Table 8 reports OLS estimates based on the baseline specification (3) in which the dependent variable is now an indicator for process innovation. We find that process innovation rises due to the EU8 immigration shock, with estimates ranging from 0.140 to 0.393. The magnitude of the estimate from the strictest specification in Column (6) indicates that the average labor market (which had just over a half percentage point increase in the share of their population from EU8 countries) was 8.4 percentage points more likely to be engaged in process innovation after the EU expansion. This is a bit larger than the actual rise in the share of firms doing process innovation of 5 percentage points (see Table 2), and comes from a baseline mean likelihood of 21 percent (see Table 2) – so is a large effect.

Notably, the effect is more than double the estimate of the reduction in product innovation, indicating that overall innovation intensity rose due to the shock. This is consistent with Table 9 which reports a specification similar to (3) but with log R&D as the dependent variable. Table 9 makes clear that innovation inputs, as reflected in R&D, indeed rose due to the shock. The low-skill labor supply shock therefore seems unambiguously to have raised innovation among UK firms, despite the (on average) shift away from product innovation.

In Appendix Table A7 we present the baseline results with London removed from the analysis, where we see that the coefficients are consistent with the baseline results.

Again, we can provide additional suggestive evidence on the baseline economic magnitudes by combining the estimates in Table 8, column (6) with the innovation premia reported in Table 1. We find that the observed average rise in the EU8 immigration share over the 2004-2008 period (of just over a half of percentage point) was associated with a $8.4 \times 0.494 = 4.14$ percent increase in firm revenue, a $8.4 \times 0.394 = 3.31$ percent increase in employment and a $8.4 \times 0.643 = 5.40$ percent increase in R&D spending.

Table 12 introduces two interactions with the treatment variable: the share of workers that are low skilled and the log of firm revenues as a proxy for firm size. The estimates indicate that process innovation was increasing in the low-skill intensity of the firm as well as in firm size. Again, consistent with the conceptual framework and the weak absolute equilibrium bias result from Acemoglu [2], an increase in the relative abundance of low-skill workers increased the intensity of low-skill biased technological investment, in this case reflected in increased process innovation by low-skill intensive firms. The estimates indicate that moving from the mean level of the low-skill share (86 percent) to all low-skill workers (100 percent) would have increased the (positive) process innovation response

by around 10 percent, a significantly smaller effect than was found for product innovation.

7 Concluding Remarks

Immigration policy is a prominent political issue— there is much debate surrounding both the optimal aggregate number of immigrants as well as the mix between skill levels of potential labor migrants. 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 find that the low-skill labor supply shock 1) increased process innovation while 2) reducing product innovation. We further find evidence in favor of a demand side role for immigrants in spurring new product creation, which was likely offset by a relative increase in the price of high-skill workers, thus discouraging net product innovation activity. Overall we find that net innovation went up, as the rise in process innovation was larger than the fall in product innovation.

More generally, the results suggest that one reason that the estimated labor market effects of immigration are small is that firms adjust their production processes rapidly in response to changes in input endowments. This is in line with Dustmann and Glitz [20] who find that German firms in the tradable sector responded to changes in available skill by altering their production technologies toward those that were more intensive in the now-relatively-more-abundant skill type. Lewis [37] provides an overview of models that provide channels of adjustment to labor supply shocks, and describes a new thread of research that uses exogenous immigration shocks, such as the one exploited here resulting from the 2004 EU expansion, to test the predictions of this class of models. Our paper suggests that indeed these types of production technology and capital adjustments seem to take place and must be taken into account in any predictions about the impact of immigrants on the relative wages of native workers.

The results with respect to product innovation are also informative. Previous work has focused almost exclusively on the direct link between high-skill immigrants and inventive activity. In this paper we explored the implications of a low-skill shock and its direct and indirect effects on product creation. A potential direction for future research is to incorporate both high- and low-skill immigration flows into a model and empirical research design that explores firm innovation choices. In this case firms may face a wider range of tradeoffs, and any empirical findings would be particularly relevant for countries for whom immigrant skill is bimodal (such as in the UK or US). A focus on longer run outcomes would also be informative for policy makers.

References

- [1] Acemoglu D. 1998. Why Do New Technologies Complement Skills? Directed Technical Change And Wage Inequality. *The Quarterly Journal of Economics* **113**: 1055–1089.
- [2] Acemoglu D. 2007. Equilibrium bias of technology. *Econometrica* **75**: 1371–1409.
- [3] 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.
- [4] Antras P, Garicano L, Rossi-Hansberg E. 2006. Organizing offshoring: Middle managers and communication costs. Technical report, National Bureau of Economic Research.
- [5] Antràs P, Rossi-Hansberg E. 2009. Organizations and trade. *Annu. Rev. Econ.* **1**: 43–64.
- [6] Basker E. 2012. Raising the barcode scanner: Technology and productivity in the retail sector. *American Economic Journal: Applied Economics* **4**: 1–27.
- [7] Bell B, Fasani F, Machin S. 2013. Crime and immigration: Evidence from large immigrant waves. *Review of Economics and statistics* **21**: 1278–1290.
- [8] Blanchflower DG, Lawton H. 2009. The impact of the recent expansion of the eu on the uk labor market. In *EU labor markets after post-enlargement migration*. Springer, 181–215.
- [9] Bloom N, Draca M, Van Reenen J. 2016. Trade induced technical change? the impact of chinese imports on innovation, it and productivity. *The Review of Economic Studies* **83**: 87–117.
- [10] Bloom N, Garicano L, Sadun R, Van Reenen J. 2014. The distinct effects of information technology and communication technology on firm organization. *Management Science* **60**: 2859–2885.
- [11] Bloom N, Sadun R, Van Reenen J. 2012. Americans do it better: Us multinationals and the productivity miracle. *The American Economic Review* **102**: 167–201.
- [12] Bodvarsson ÖB, Van den Berg HF, Lewer JJ. 2008. Measuring immigration’s effects on labor demand: A reexamination of the mariel boatlift. *Labour Economics* **15**: 560–574.
- [13] Card D. 2012. Comment: The elusive search for negative wage impacts of immigration. *Journal of the European Economic Association* **10**: 211–215.
- [14] Cortes P. 2008. The effect of low-skilled immigration on us prices: evidence from cpi data. *Journal of political Economy* **116**: 381–422.
- [15] Dhingra S. 2013. Trading away wide brands for cheap brands. *The American Economic Review* **103**: 2554–2584.

- [16] Di Giovanni J, Levchenko AA, Ortega F. 2015. A global view of cross-border migration. *Journal of the European Economic Association* **13**: 168–202.
- [17] Duranton G, Puga D. 2001. Nursery cities: Urban diversity, process innovation, and the life cycle of products. *American Economic Review* **91**: 1454–1477.
- [18] Dustmann C, Fabbri F, Preston I, Wadsworth J. 2003. Labour market performance of immigrants in the uk labour market. *Home Office Online Report 05/03* .
- [19] Dustmann C, Frattini T, Halls C. 2010. Assessing the fiscal costs and benefits of a8 migration to the uk. *Fiscal Studies* **31**: 1–41.
- [20] Dustmann C, Glitz A. 2015. How do industries and firms respond to changes in local labor supply? *Journal of Labor Economics* **33**: 711–750.
- [21] Eckel C, Neary JP. 2010. Multi-product firms and flexible manufacturing in the global economy. *The Review of Economic Studies* **77**: 188–217.
- [22] Eizenberg A. 2014. Upstream innovation and product variety in the us home pc market. *Review of Economic Studies* **81**: 1003–1045.
- [23] Elsner B. 2013. Emigration and wages: The eu enlargement experiment. *Journal of International Economics* **91**: 154–163.
- [24] Feenstra R, Ma H. 2007. Optimal choice of product scope for multiproduct firms under monopolistic competition. Technical report, National Bureau of Economic Research.
- [25] Feenstra RC. 1994. New product varieties and the measurement of international prices. *The American Economic Review* **84**: 157–177.
- [26] Fogel M, Peri G. 2016. Immigrants’ effect on native workers: New analysis on longitudinal data. *American Economic Journal: Applied Economics* **8**: 1–34.
- [27] Gaggl P, Wright GC. 2017. A short-run view of what computers do: Evidence from a uk tax incentive. *American Economic Journal: Applied Economics* **9**: 262–94.
- [28] Gray R, Montresor G, Wright GC. 2020. Replication files for processing immigration shocks: Firm responses on the innovation margin.
URL [10.17632/w543df97pr.1](https://doi.org/10.17632/w543df97pr.1)
- [29] Hanson GH, Slaughter MJ. 1999. The rybczynski theorem, factor-price equalization, and immigration: evidence from us states. Technical report, National Bureau of Economic Research.
- [30] Holmes C. 2015. *John Bull’s island: immigration and British society, 1871-1971*. Routledge.
- [31] 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.

- [32] Huergo E, Jaumandreu J. 2004. Firms' age, process innovation and productivity growth. *International Journal of Industrial Organization* **22**: 541–559.
- [33] Hunt J. 2011. Which immigrants are most innovative and entrepreneurial? distinctions by entry visa. *Journal of Labor Economics* **29**: 417–457.
- [34] Kerr WR. 2013. Us high-skilled immigration, innovation, and entrepreneurship: Empirical approaches and evidence. Technical report, National Bureau of Economic Research.
- [35] Lafortune J, Tessada J, Lewis E. 2019. People and machines: A look at the evolving relationship between capital and skill in manufacturing 1860-1930 using immigration shocks. Technical Report 1.
- [36] Lewis E. 2011. Immigration, skill mix, and capital skill complementarity. *The Quarterly Journal of Economics* **126**: 1029–1069.
- [37] Lewis E. 2013. Immigration and production technology. *Annu. Rev. Econ.* **5**: 165–191.
- [38] Maré DC, Fabling R, Stillman S. 2014. Innovation and the local workforce. *Papers in Regional Science* **93**: 183–201.
- [39] Markus ML, Robey D. 1988. Information technology and organizational change: causal structure in theory and research. *Management science* **34**: 583–598.
- [40] Mazzolari F, Neumark D. 2012. Immigration and product diversity. *Journal of Population Economics* **25**: 1107–1137.
- [41] Ottaviano GI, Peri G. 2012. Rethinking the effect of immigration on wages. *Journal of the European economic association* **10**: 152–197.
- [42] Ozgen C, Nijkamp P, Poot J. 2013. Measuring cultural diversity and its impact on innovation: Longitudinal evidence from dutch firms. *IZA Discussion Paper No. 7129* .
- [43] Peri G, Sparber C. 2009. Task specialization, immigration, and wages. *American Economic Journal: Applied Economics* **1**: 135–169.
- [44] 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.

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.