

In Search of Labor Demand[†]

By PAUL BEAUDRY, DAVID A. GREEN, AND BEN M. SAND*

We propose and estimate a novel specification of labor demand which encompasses search frictions and the role of entrepreneurs in new firm creation. Using city-industry variation over four decades, we estimate the wage elasticity of employment demand to be close to -1 at the industry-city level and -0.3 at the city level. We argue that the difference between these estimates reflects the congestion externalities predicted by the search literature. Our estimates also indicate that entrepreneurship should be treated as a scarce factor in the determination of labor demand. We use our estimates to evaluate the impact of large changes in the minimum wage on employment. (JEL J23, J31, J38, L26, M13, R23)

North American policymakers interested in how wage costs affect employment decisions could be excused for being confused by what the economics literature has to tell them. At one extreme, studies using variation in minimum wages and payroll taxes tend to find only small wage elasticities of employment demand (Blau and Kahn 1999). At the other, studies of regional responses to labor supply shocks generally find small wage impacts and large employment changes, which is suggestive of very elastic labor demand (Blanchard and Katz 1992; Krueger and Pischke 1997).¹ Further variation in estimates arises in the literature because different studies use different units of observation, time frames, and identification strategies, often without a clear reference to theory to support their choice. Our goal in this paper is to propose and estimate a new specification for labor demand that is based on a comprehensive view of the labor market and that is capable of reconciling different findings in the literature.

* Beaudry: Vancouver School of Economics 6000 Iona Drive, Vancouver, Canada, V6T 1L4, and NBER (email: Paul.Beaudry@ubc.ca); Green: Vancouver School of Economics 6000 Iona Drive, Vancouver, Canada, V6T 1L4, and IFS London (email: david.green@ubc.ca); Sand: York University, 1090 Vari Hall, 4700 Keele Street, Toronto, ON, M3J 1P3 (email: bmsand@yorku.ca). This paper was accepted to the *AER* under the guidance of Luigi Pistaferri, Coeditor. We thank Chris Bidner, Bruno Deceuse, David Card, Giovanni Gallipoli, Patrick Kline, Alex Lefter, Kevin Milligan, and Richard Rogerson for comments and Vadim Marmer for econometric advice. This paper incorporates some results previously circulated under the title “The Elasticity of Job Creation.”

[†] Go to <https://doi.org/10.1257/aer.20141374> to visit the article page for additional materials and author disclosure statement(s).

¹ The two extremes are captured in the minimum wage literature on one end (where studies commonly find either small positive or small negative elasticities) and the literature on city adjustments to immigration shocks on the other (where, for example, Card 1990a finds virtually no wage response to the Mariel Boatlift supply shock in Miami). The small elasticities related to wage policy shifts, though, may not apply to continental Europe. Both Kramarz and Philippon (2001) and Cahuc, Carcillo, and Barbanchon (2014) find substantial elasticities for France, for example. Uncovering the reasons for differences between Europe and North America would certainly be interesting but is beyond the scope of this paper. We focus on estimates using US variation.

A natural starting place to look for answers regarding the wage elasticity of employment is the micro-literature on firm demand for labor (see, for example, Hamermesh 1993, ch. 4). The goal of this literature has traditionally been to estimate how the average firm responds to a change in wages, generally holding total output constant. It is a literature that is very close in spirit to the literature estimating production functions. Knowing the properties of a firm's production functions, such as the extent of capital labor substitutability, is certainly interesting. However, it is unlikely to provide a complete assessment of how total labor demand within a market responds to a change in wages. For example, a production function perspective on labor demand will necessarily miss any adjustment on the extensive margin since entry and exit decisions of firms are excluded. Moreover, when discussing responses at the market level, it is not particularly interesting to keep the output produced by firms fixed.

A firm perspective on labor demand may also differ from a market perspective because of search and matching frictions. When adopting a firm perspective, a change in the wage is viewed as affecting the firm's employment decision, and any external effects of that employment decision on the decisions of other firms is not considered. However, in the presence of search and matching frictions, an increase in the employment of one firm has a direct externality effect on the employment decisions of other firms, even holding wages fixed, since it increases market tightness and thereby increases the cost of recruitment. Such a mechanism may imply a difference between the market response to a change in labor costs and the simple sum of isolated firm responses. If one is interested in how labor demand in a market responds to wages, one must move away from a perspective focused at the individual firm level and adopt an approach that explicitly takes into account the many channels through which changes in wage costs can affect employment decisions. Accordingly, our approach will be to derive an empirically tractable specification of the market-level demand for labor that takes into account several different margins of adjustment.²

The labor demand specification we propose is built from micro-foundations and incorporates four main determinants of employment. The first is a direct wage effect of the kind that is central to any study of labor demand. In our framework, this effect will capture adjustments on both the intensive and entry margins of firm decisions. Second, there is a labor market tightness effect aimed at capturing the congestion externalities emphasized in the search and matching literature. Third, we also include population size as a determinant of employment demand. From the perspective of the traditional labor demand literature this is unconventional because one would typically expect population size to determine labor supply, not labor demand. However, once one models the process of firm creation explicitly, and recognizes that entrepreneurs may be a limiting factor in job creation, it becomes necessary to include population size as a determinant of employment demand since it reflects the extent of entrepreneurship opportunities. Finally, there are the effects of technological change that will appear in the error term of our specification.

²Our focus on medium-run wage effects on employment differentiates our work from studies of regional adjustment to aggregate labor demand changes (Blanchard and Katz 1992; Bartik 1993, 2009) which mainly focus on unemployment dynamics.

In the empirical section of the paper, we estimate our labor demand specification at both the industry-city level and at the aggregate city level using data from the 1970, 1980, 1990, and 2000 US Censuses and the American Community Survey for the pairs of years 2007–2008 and 2014–2015. Our approach is to treat the cities as observations on a set of local economies, allowing us to identify within-city general equilibrium effects. Since we look at changes in employment outcomes over (roughly) ten-year periods, our focus will clearly be on medium-run adjustment, and, for this reason, our approach will downplay certain adjustment costs that have been central to the dynamic labor demand literature that generally focuses on much higher frequency decisions (for example, Cooper, Haltiwanger, and Willis 2015; King and Thomas 2006; Kramarz and Michaud 2010).

As is common in all studies of demand or supply, the key difficulty is finding convincing data variation that allows consistent estimation of the causal impacts of the variables of interest. To this end, we rely on a set of instruments that are similar in spirit to that first proposed by Bartik (1991) to identify each of our key labor demand determinants. The instruments we build use developments at the national level to predict local outcomes and rely on the identifying assumption that changes in productivity at the local level are independent of past levels of local productivity. We discuss the plausibility of this assumption, which is certainly questionable, and provide a very informative over-identification test. To identify wage effects, we build two instruments (discussed in more detail below) that are based on our earlier work on search models in a multisector context (Beaudry, Green, and Sand 2012—henceforth, BGS). To identify the labor market tightness effect, we exploit the commonly used Bartik instrument.³ Finally, to identify an effect of population size on labor demand, we use a variant of the commonly used ethnic enclave instrument from the immigration literature (which is also a Bartik-style instrument).

Since our main focus in this paper is on consistently estimating the wage elasticity of labor demand, it is worth providing some extra detail on our identification strategy for these wage effects up front. In our earlier work (BGS), we argue that wage patterns in the United States indicate that wages are at least partially the outcome of a bargaining process that takes place at the industry-city level. In that process, the outside option of workers is an important determinant of the wage. In BGS we point out that the outside option for workers in a particular industry-city cell is better if the industrial composition of employment in the city is weighted toward high-paying industries. That is, a worker in, say, construction can bargain a better wage if the city he lives in includes a high-paying steel plant instead of a lower-paying textile plant since one of his outside options is to move to the steel mill. BGS shows how to build instruments for wage changes that are based on this insight.⁴ These instruments are

³This instrument was first presented by Bartik (1991) and has been used in much subsequent work (for example, Bartik 1993; Blanchard and Katz 1992; Bound and Holzer 2000). The Bartik instrument corresponds to a prediction of employment growth in a city based on industrial growth rates at the national level combined with start-of-period employment composition in the city.

⁴The idea of obtaining identification using variation in workers' outside options has precedents in the literature examining union wage and employment contracts (e.g., Brown and Ashenfelter 1986; MaCurdy and Pencavel 1986; Card 1990b) as these papers exploit measures of alternative wages outside the specific contract in their estimation. Card (1990b) finds that the real wage in manufacturing has a positive effect on wage changes in the Canadian union contracts he studies, which echoes the mechanism underlying our basic source of identification. In a similar spirit, MaCurdy and Pencavel (1986) obtain estimates of production function parameters from data on wage and

of a similar form to the classic Bartik instrument in the sense that they rely on an assumption that productivity growth in a city is not related to the initial employment composition in the city. Since we can build more than one instrument based on this outside option insight, we can use an over-identifying test to evaluate the plausibility of the underlying identification restrictions. We show that this test is quite strong and that it is passed easily in our data.

The main empirical results of the paper are as follows. We find a statistically significant and economically meaningful negative trade-off between city-level employment rates and wages over ten-year periods. When looking at the industry level within a city, we find that a 1 percent increase in the wage in an industry-city cell leads to a decrease in the employment rate in that cell of approximately 1 percent. This result holds both when we look at all industries and when we look only at industries producing highly traded goods. When looking at the city level, we find that a 1 percent increase in the wages within all industries in a city leads to only a 0.28 percent decrease in the employment rate.⁵ We argue that the smaller effect at the city level compared to the industry level reflects the impact of search externalities. In particular, we interpret the latter result as reflecting that a wage increase in all industries leads to a less tight labor market, thereby reducing search costs to firms. The fall in search costs partially compensates for the increase in wage costs, leading to a smaller fall in employment than would have happened if wages only increased in a worker's own industry.⁶ Finally, we find that an increase in population holding wages constant leads to an approximately proportional increase in labor demand.⁷ We interpret this finding as indicating that the number of entrepreneurs involved in creating jobs in a city moves proportionally with the size of the city. Moreover, we will argue that this population size result also indicates that local labor markets are unlikely to be significantly constrained by fixed physical factors when looking over a ten-year period.

An important implication of our findings relates to the identification of wage cost effects. In particular, our results imply that shifts in population caused by migration shocks cannot be used as instruments for the wage in labor demand specifications because population size is a direct determinant of labor demand. Put a different way, the wage-employment trade-offs estimated in the literature using immigration shocks do not correspond to the wage elasticity of labor demand that is of concern to most policymakers. In our view, the relevant wage elasticity of labor demand for many policy issues needs to be estimated holding population size constant.⁸

employment setting for typesetters when allowing for an alternative wage to affect the efficient outcome through an impact on union preferences.

⁵In Hamermesh (1993), the main estimates he reports lie in a range near -0.3 , which suggest a rather low elasticity of substitution between capital and labor. While this elasticity is numerically very close to the one we obtain here, it is not appropriate to compare them as they do not address the same question.

⁶Our finding of smaller effects at the city versus the industry level suggests that any possible positive demand linkages across industries in a city are dominated by the negative search externalities.

⁷One implication from this is that specifications with the employment rate rather than the employment level as the dependent variable are appropriate. Our reading of the existing labor demand literature is that papers use either employment levels or employment rates without providing any direct rationale for their decision.

⁸It is interesting to think of this result in the context of the employment effects estimated in, for example, Card's (1990a) work on the effects of the Mariel Boatlift. Card shows that the sizable inflow of Cuban refugees into the Miami labor market had little effect on wages. In the context of our extended model, if the inflow of migrants brings with it a proportional number of entrepreneurs then one should observe something like a replication of the existing economy; that is, a one-for-one increase in employment with little change in wages. However, according to our

The crux of our findings is found in the combination of a modest negative wage elasticity and the result that, holding wages fixed, increases in labor supply increase employment one-for-one. We believe that these findings are easiest to interpret in terms of models with explicit recognition of entrepreneurs. In particular, within our framework these results imply that (i) entrepreneurs face a span of control problem or at least downward-sloping demand for their product and (ii) the elasticity of the supply of entrepreneurial talent to higher profits is far from perfectly elastic. Our estimates suggest that both these elements have to be present to explain the data. Overall, we view our results as highly supportive of labor market models that emphasize the role of scarce entrepreneurial talent in the job creation process. In keeping with this, in a system estimation in which we break wage effects down into an intensive margin (i.e., changes in employment holding the number of entrepreneurs constant) and an extensive margin (i.e., changes in the number of entrepreneurs), the wage effect operates almost exclusively at the extensive margin.

In the last section of the paper, we use our estimates of labor demand to evaluate the effects of a large increase in the minimum wage on employment. Recently, there have been several cities in the United States that have either substantially increased the legal minimum wage or are expected to do soon. For example, the minimum wage in the metropolitan area of Seattle is scheduled to increase from \$9.41 to \$15 in the coming years. Our model is well suited to examine the long-run response of the local employment rate to such a change. Our estimates imply that an increase in the minimum wage of the size being implemented in Seattle will lead to an overall decrease in the employment rate of 2.1 percentage points. Given our estimates of the different margins of adjustment, the short-run impact may be quite small, but over time the impact should grow due to reduced entrepreneurial activity.

The remaining sections of the paper are structured as follows. In Section I, we derive our empirical specifications for labor demand. We begin deriving a labor demand specification assuming that employers can readily hire workers at the going wage. We then extend our approach to allow for search frictions and emphasize how greater tightness in the labor market should negatively affect employment at the industry level. In Section II, we discuss issues related to identification of parameters. In particular, we present and justify the instrumental variable (IV) strategy we exploit for estimation. In Section III, we discuss the data and our construction of variables. In Section IV, we report our main empirical results. In Section V, we examine the robustness of our results to breakdowns by education and to incorporating slow adjustment of labor and wages. Section VI contains a summary of the main empirical results and our interpretation of them. Finally, Section VII uses estimates to evaluate the effect of a large increase in the minimum wage and Section VIII offers concluding remarks.

I. Deriving Labor Demand

Our goal in this section is to derive an empirically tractable specification for the locus describing the trade-off between wages and employment demand at the level of

work, this should not be interpreted as implying a perfectly elastic labor demand curve. It simply reflects the fact that holding wages constant, employment tends to increase with the size of population.

an industry or a whole economy. While it may seem natural to refer to that locus as a labor demand curve (and we will describe it in those terms as we proceed), there is a sense in which this terminology can be misleading. In particular, the traditional labor demand literature has focused on identifying parameters of production functions that are relevant for firm-level employment decisions. While our approach will include such elements, we will also allow for effects of elements related to the entry process of firms and to search frictions, as both of these can affect the policy-relevant trade-off that is of interest to us. As we will see, if those elements are relevant then they imply that what we will want to estimate is an equilibrium locus that reflects features beyond what are captured in the labor demand curve of any one firm.

In order to make the derivation of the labor demand locus transparent, it is helpful to use a metaphor in which firms hire workers through competitive recruitment agencies, and it is only the recruitment agencies that face search frictions when looking for workers. This allows us to temporarily disregard search frictions and first focus on the determination of firm employment and entry decisions in industry i in city c , taking wages, w_{ic} , set by the recruitment agencies as given. To this end, consider an environment where the good produced in industry i is traded on a national market at a price p_i , and where physical capital can be rented on the national market at a rental price r . We will also allow for the possibility of a local fixed factor in each industry, which can be rented at price p_{ic}^X . Each potential entrepreneur in this market has access to a production function $F^i(e_{ic}^j, K_{ic}^j, X_{ic}^j, \theta_{ic})$, where e_{ic}^j is the number of workers employed by entrepreneur j in industry i in city c , K_{ic}^j is capital rented by the entrepreneur, X_{ic}^j is the use of the locally fixed factor, and θ_{ic} is an exogenous productivity parameter capturing comparative advantage in the industry-city cell. We assume, for the moment, that there is only one type of labor. We discuss how to extend the framework to take into account worker heterogeneity in Section IIA. To ease presentation, we will assume that the production function takes the Cobb-Douglas form $F^i(e_{ic}^j, K_{ic}^j, X_{ic}^j, \theta_{ic}) = (e_{ic}^j)^{\alpha_1} (K_{ic}^j)^{\alpha_2} (X_{ic}^j)^{\alpha_3} \theta_{ic}$, with $0 < \alpha_1 + \alpha_2 + \alpha_3 \leq 1$. We will point out where restrictions imposed by the Cobb-Douglas form affect our conclusions and describe how they are extended by relaxing that assumption.

If entrepreneur j decides to enter the market, optimization implies that the employment level at his firm will be given by

$$(1) \quad e_{ic}^j = \left[r^{-\alpha_2} p_{ic}^{-\alpha_3} \alpha_1^{(1-\alpha_2-\alpha_3)} \alpha_2^{\alpha_2} \alpha_3^{\alpha_3} \right]^{\frac{1}{1-\alpha_1-\alpha_2-\alpha_3}} (w_{ic})^{\frac{-(1-\alpha_2-\alpha_3)}{1-\alpha_1-\alpha_2-\alpha_3}} (\theta_{ic} p_i)^{\frac{1}{1-\alpha_1-\alpha_2-\alpha_3}}.$$

The issue that interests us is how to go from this firm-level labor demand to aggregate labor demand in industry i in city c . The answer to this question depends on how we specify the firm's entry process, whether we assume the presence of a span of control problem⁹ and how we model the market for locally fixed factors. If there is no span of control problem, then going from firm demand to market demand is trivial since firm size is indeterminate and therefore the firm and the market are interchangeable. This is an approach sometimes taken in the labor demand literature.

⁹In the context of this production function, span of control problems are captured by assuming that there are decreasing returns to scale at the firm level, that is, $\alpha_1 + \alpha_2 + \alpha_3 < 1$.

In contrast, we will allow for a span of control problems at the firm level, implying that firm size is not generally indeterminate and that entry decisions become relevant. Accordingly, we need to discuss how to capture tractably the effects of firm entry on labor demand.

When considering firm entry, there are two simple but extreme cases. At one extreme, we could follow the firm entry literature (e.g., Hopenhayn 1992) and assume that there is an infinite supply of potential entrants, with each entrant needing to pay a common fixed cost upon entry. We see this situation as extreme since it leads to a labor demand curve which has an infinite elasticity with respect to wages (holding other prices fixed), thus presupposing an answer to the question we are investigating. At another extreme is the assumption that the supply of entrepreneurs is determined exogenously. This assumption is also unattractive since it does not allow for wages or profits to feedback on entry decisions. Our objective is to choose a specification of the entry process which allows for these extremes but does not impose them.¹⁰ We do this by starting with a set of potential entrepreneurs who could operate firms in industry i , N_i . We assume that these people can live in any city and receive opportunities in their own city or any other. A potential entrepreneur becomes aware of an opportunity in a particular city with probability ψ_{ic} . We assume that ψ_{ic} is proportional to the size of the city, so that the number of potential entrepreneurs who are aware of options in a given industry-city cell is given by $N_{ic} = \gamma_{0i} L_c^{\gamma_1}$, where L_c is the size of the local market as given by its labor force, γ_{0i} is an industry specific effect that is constant across markets and incorporates N_i , and $\gamma_1 \geq 0$. The γ_1 parameter determines whether the set of knowledgeable potential entrepreneurs is proportional ($\gamma_1 = 1$), less than proportional ($\gamma_1 < 1$), or more than proportional ($\gamma_1 > 1$) in size to the local population, and its value is an empirical question.

We allow for heterogeneity among entrepreneurs by incorporating a margin at which potential entrepreneurs do or do not become active, in the sense of opening a firm and hiring workers. We assume that when a potential entrepreneur learns about an opportunity in a particular city \times industry cell, this involves learning about the profits in that cell (which are common to all firms in the cell) as well as about the entrepreneur's own cost of opening a firm there. The latter cost, f , is drawn from a distribution with cumulative distribution function $G(f)$. The heterogeneity among entrepreneurs leads to a cutoff rule where only potential entrepreneurs with a fixed cost below some cutoff f^* (equal to the profits from operating in that industry and city) will become active in a market. To allow for simple analytic expressions, we further assume that $G(f)$ takes the form $G(f) = (f/\Gamma)^\phi$, where $0 \leq \phi$ and $f \in [0, \Gamma]$. With this formulation of the distribution of the entry costs, the extreme case where there is only one common fixed cost at Γ can be captured in the limit when ϕ goes to infinity. Thus, the number of active entrepreneurs in an industry \times city cell is the product of the number of potential entrepreneurs in that cell times the fraction of those entrepreneurs who find it profitable to operate a business, i.e., $N_{ic}^a = N_{ic} G(f^*) = \gamma_{0i} L_c^{\gamma_1} G(f^*)$.¹¹

¹⁰ Kuhn (1988) similarly considers wage and employment setting when the number of entrepreneurs is endogenously, but not perfectly elastically, determined. His focus is on the interaction of unionization and entrepreneurship.

¹¹ We can incorporate the entrepreneurs choosing between opening a firm in the city where they get a business option and remaining as a paid worker in their current city. If potential entrepreneurs are spread around the country, this would introduce a national average wage into our specification that would become part of the general time

Total employment demand in industry i in city c , which we will denote by E_{ic} , is the result of the aggregation of employment across all active entrepreneurs as given by $E_{ic} = N_{ic}^a \cdot e_{ic}$. Thus, total employment in an industry-city cell is given by

$$(2) \quad E_{ic} = \gamma_{0i} e_{ic} G(f^*) L_c^{\gamma_1}.$$

In the absence of any fixed factors ($\alpha_3 = 0$), it is straightforward to replace e_{ic} and f^* in (2) to express total employment in an industry-city cell as a function of wages. In the presence of a fixed factor, one needs to also take into account that the price of the fixed factor will adjust to clear its market. Taking all these effects into account, we get the following expression in log form for employment in industry i in city c :

$$(3) \quad \ln E_{ic} = \alpha_{0i} - \frac{1 - \alpha_2 + \phi(\alpha_1 + \alpha_3)}{1 - \alpha_1 - \alpha_2 + \alpha_3 \phi} \ln w_{ic} + \frac{\gamma_1(1 - \alpha_1 - \alpha_2 - \alpha_3)}{1 - \alpha_1 - \alpha_2 + \alpha_3 \phi} \ln L_c + \epsilon_{ic},$$

where α_{0i} captures industry effects (such as price of the good and the rental rate of capital) which are common across cities and the error term is $\epsilon_{ic} = \frac{1 + \phi}{1 - \alpha_1 - \alpha_2 + \alpha_3 \phi} (\ln \theta_{ic} + \alpha_3 X_{ic})$.

A first aspect to note about equation (3) is the appearance of local population size on the right-hand side. This reflects the fact that the set of potential entrepreneurs who are aware of opportunities in the city is taken to be increasing with population size. We believe it is important to allow for a potential role for population size in the determination of labor demand, in part, because of its implications for identifying the wage elasticity of labor demand. In particular, if population size is important then how wages adjust in response to a change in the population in a local labor market (e.g., due to an immigration shock) may reveal nothing about the wage elasticity of labor demand implied by (3). An increase in population in a local labor market could be met with a proportional increase in employment at fixed wages. In a standard labor demand specification, this would be interpreted as implying that the labor demand curve is perfectly elastic while in the specification here, this pattern could be observed even if the wage elasticity of labor demand in (3) is very small. Allowing for population size to affect labor demand is also useful because looking at how population growth affects employment holding wages fixed can provide substantial information about the functioning of the labor market. For example, if population enters into this equation with a coefficient of 1 then one potentially can infer that entrepreneurship is proportional to the population and that fixed factors are unimportant. In such a case, it would be more appropriate to describe the wage-employment trade-off in a market as one between wages and the employment rate as opposed to one between wages and the level of employment. Our empirical results

term in our empirical specification. We omit this element for simplicity. Similarly, we could allow for the arrival of opportunities in an industry \times city cell to be proportional to the productivity of that cell, θ_{ic} , as well as city size. That would alter the coefficient on $\ln \theta_{ic}$ in the error term but otherwise leave our specification unchanged and so we also omit this for simplicity.

will, in fact, support a wage-employment rate specification as we will show that employment appears to increase one-for-one with population, holding wages fixed.

We next turn our focus to the coefficient on the wage in (3). This coefficient is always negative since it is given by $-\frac{1 - \alpha_2 + \phi(\alpha_1 + \alpha_3)}{1 - \alpha_1 - \alpha_2 + \phi\alpha_3}$. There are two scenarios under which this coefficient equals minus infinity, i.e., where there is perfectly elastic demand.¹² Both require that there be no fixed factor ($\alpha_3 = 0$). First, if there is no span of control problem (and there are no fixed factors), then $1 - \alpha_1 - \alpha_2 = 0$ and the wage elasticity becomes infinite. Alternatively, if potential entrepreneurs all face the same cost of entry, Γ , then ϕ must equal infinity as there is a mass point in the function $G(\cdot)$, which again implies an infinitely elastic labor demand curve. Importantly, for the wage elasticity to be less than infinite, neither of these conditions can hold. Hence, finding evidence of both a less than infinite wage elasticity and that employment moves proportionally to population conditional on the wage can be interpreted as indicating the likely presence of both a span of control problem and that there is not an infinitely elastic supply of entrepreneurs waiting to take advantage of any profit opportunity.¹³

We see the introduction of a fixed physical factor in our setup as helping clarify the empirical differences implied by a span of control problem versus the presence of a fixed factor. With the presence of a fixed factor, the effect of population on labor demand is likely to be smaller than 1 even if $\gamma_1 = 1$; that is, even if available entrepreneurs are proportional to the population. This is intuitive as population growth will cause the fixed factor to become more constraining even in the presence of more entrepreneurs. However, and most importantly, if we assume away the span of control problem ($1 - \alpha_1 - \alpha_2 - \alpha_3 = 0$), then even if $\gamma_1 > 0$, population will not enter the labor demand specification. The presence of a fixed factor can justify why the wage elasticity of labor demand may be less than minus infinity but it cannot rationalize why an increase in population may be met with increased employment at fixed wages. To rationalize this, while maintaining the feature that the wage elasticity of labor demand is less than infinite, one needs the presence of a limiting factor that grows with population. Entrepreneurs play that role in our framework.

A. Recruitment Agencies and Search Frictions

In our derivation of equation (3), we assumed that firms did not face search frictions in the labor market as they could hire workers from recruitment agencies at the going wage. In this subsection, we examine how the cost of labor to firms changes

¹²We derived equation (3) under the assumption that all goods in an industry are perfect substitutes. If, instead, we assume that goods from each entrepreneur are a differentiated product then there is a further reason, beyond the span of control problem, for a fall in wages to have a limited effect on employment demand within a firm. Since the extension of the specification above to the case where the outputs of the different entrepreneurs are not perfect substitutes is rather straightforward, we omit it here. However, it should be noted that such an extension does change the interpretation of the coefficient on wages from one that is driven mainly by the span of control problem and firm entry decisions to one that also takes into account the substitutability of products within the industry.

¹³A more subtle issue in equation (3) is the implicit restriction on the wage elasticity of labor demand embedded in it. For example, if $\alpha_3 = 0$ this elasticity should always be greater than 1 in absolute value. This feature is actually an artifact of the Cobb-Douglas structure and does not hold for more general production functions. For this reason, it should not be viewed as a relevant restriction.

if the recruitment agencies face search frictions. We continue to assume that entrepreneurial firms rent labor from recruitment firms at a price w_{ic}^z , which they take as given. This allows us to introduce search frictions while leaving the other elements of firm behavior (deciding whether to produce and how much to buy of the different inputs if production takes place) the same as in the previous section.

In order to introduce search frictions faced by the recruitment firms, we need to be explicit about the dynamic setting. Accordingly, we will assume that time is continuous and that all the costs facing entrepreneurs discussed previously now represent instantaneous costs for flow services. Since the decision faced by our entrepreneurial firms did not involve a dynamic component, the demand for labor as given by (3) is valid at every point in time. We now assume the existence of a large set of recruitment agencies, each of which can decide whether to post a vacancy at any point in time, where a vacancy is aimed to hire someone for a specific industry. The flow cost of posting a vacancy for industry i in city c is denoted h_{ic} . When a recruitment agency finds a worker, it offers the labor services on the market and obtains a flow return of $w_{ic}^z - w_{ic}$, where w_{ic} is now the wage received by the worker which in general is less than the wage, w_{ic}^z , paid by the good producing firms. Workers are hired from a common pool, regardless of the industry to which they will ultimately be allocated. Job vacancies and unemployed workers match according to a constant returns to scale matching function given by $M(L_c - E_c, V_c) = (L_c - E_c)^\nu V_c^{1-\nu}$, where E_c is total employment in the city and V_c is the number of vacancies. Given this matching function, the flow rate at which a recruitment agency finds a worker is given by $\left(\frac{L_c - E_c}{V_c}\right)^\nu$. Assuming that matches break up exogenously at rate δ , the steady-state flow rate at which recruitment agencies find workers will be given by $\left[\frac{1}{\delta} \left(\frac{1}{E_c/L_c} - 1\right)\right]^{\frac{\nu}{1-\nu}}$. Letting ρ denote the discount rate for these firms, the equilibrium condition imposing that the value of a vacancy be zero implies the following simple expression between w_{ic}^z and w_{ic} .¹⁴

$$(4) \quad w_{ic}^z = w_{ic} + \frac{(\rho + \delta) h_{ic}}{\left[\frac{1}{\delta} \left(\frac{1}{E_c/L_c} - 1\right)\right]^{\frac{\nu}{1-\nu}}}.$$

In (4), we see that the rental price of labor paid by firms is equal to the wage paid by the recruitment agency to workers plus a term capturing the cost of search. If h_{ic} were equal to 0, there would be no search costs and therefore w_{ic}^z would simply be equal to the wage. The importance of this search cost for the rental cost of labor depends on how firms discount the future, on the job destruction rate, and, most importantly, on the average time a recruitment agency spends searching for a worker

¹⁴To derive this relationship, we use the fact that the value of a filled job for a recruitment agency, which we can denote by J , must satisfy $\rho J = w_{ic}^z - w_{ic} + \delta(W - J)$ where W is the value of a vacancy. We combine this with the fact the W must satisfy $\rho W = -h_{ic} + \left[\frac{1}{\delta} \left(\frac{1}{E_c/L_c} - 1\right)\right]^{\frac{\nu}{1-\nu}}(W - J)$, and $W = 0$.

which is given by $1 \left[\frac{1}{\delta} \left(\frac{1}{E_c/L_c} - 1 \right) \right]^{\frac{\nu}{1-\nu}}$.¹⁵ In this latter expression, it is important to note that time spent looking for a worker can be expressed as an increasing function of the employment rate in the city: the tighter is the labor market, the higher is the employment rate and the longer it takes to fill a vacancy. Hence, the flow cost of the labor service, w_{ic}^z , will be greater in a tighter labor market, holding wages fixed. If we simplify matters further by assuming that the cost of posting a vacancy, h_{ic} , is proportional to the wage in the industry-city cell (that is, $h_{ic} = h_i \cdot w_{ic}$), then we can use (4) and (3) to get the following generalized demand for labor relationship, which now includes a term that reflects search frictions:

$$(5) \quad \ln E_{ic} = \beta_{0i} + \beta_1 \ln w_{ic} + \beta_1 \ln \left(1 + \frac{(\rho + \delta) h_i}{\left[\frac{1}{\delta} \left(\frac{1}{E_c/L_c} - 1 \right) \right]^{\frac{\nu}{1-\nu}}} \right) + \beta_2 \ln L_c + \epsilon_{ic},$$

where $\beta_1 = -\frac{1 - \alpha_2 + \phi(\alpha_1 + \alpha_3)}{1 - \alpha_1 - \alpha_2 + \alpha_3 \phi}$, $\beta_2 = \frac{\gamma_1(1 - \alpha_1 - \alpha_2 - \alpha_3)}{1 - \alpha_1 - \alpha_2 + \alpha_3 \phi}$, $\epsilon_{ic} = \frac{1 + \phi}{1 - \alpha_1 - \alpha_2 + \alpha_3 \phi} \times (\ln \theta_{ic} + \alpha_3 X_{ic})$, and β_{0i} again captures industry-specific terms.

Equation (5) provides, in our view, a simple but rich framework for exploring the trade-off between wages and employment demand. In particular, this specification departs from traditional labor demand specifications by embedding elements of both the search and firm entry with span of control literatures. As a result, our specification for labor demand includes a wage effect, a search cost effect, and a population effect: the latter two not being commonly included in traditional specifications of labor demand. Note that the coefficient on population, β_2 , will equal 1 if $\gamma_1 = 1$ and $\alpha_3 = 0$; that is, when available entrepreneurs are proportional to the population and there is no fixed physical factor. This is an important special case and one that, we will see, appears to be supported by the data. It is relevant to recall that we used steady-state conditions for the search process to derive this equation. Thus, (5) is most likely appropriate for studying medium-run or longer outcomes, and this is what we will do in our empirical work.¹⁶

In our empirical work, we will actually focus on a log-linear approximation of this equation so as to emphasize the first-order effects of the aggregate employment rate, E_c/L_c , on the industry-specific employment rate, E_{ic}/L_c . In particular, we will generally work with the equation in the form

$$(6) \quad \Delta \ln E_{ic} \approx \Delta \beta_{0i} + \beta_1 \Delta \ln w_{ic} + \beta_3 \Delta \ln \frac{E_c}{L_c} + \beta_2 \Delta \ln L_c + \Delta \epsilon_{ic},$$

¹⁵ In the search literature, it is most common to use the ratio of vacancies to unemployed workers as the measure of tightness. However, at the steady state, the unemployment to vacancy ratio can be written as a simple function of

the employment rate. In particular with the matching function in Cobb-Douglas form, $\frac{L-E}{V} = \frac{(1-\frac{E}{L})^{\frac{1}{1-\nu}}}{(\delta \frac{E}{L})^{\frac{1}{1-\nu}}}$.

¹⁶ Out of steady state, the link between price of labor w_{ic}^z and the wage paid to workers would be more complicated than what is given in (4), as the search cost could not be summarized by a function of the current employment rate.

where β_1 and β_2 are unchanged from before. The term β_3 can be written as $\beta_3 = \beta_1 \Phi$, where Φ is greater than 0 and corresponds to the log-linear approximation of $1 + \frac{(\rho + \delta)h_i}{\left[\frac{1}{\delta}\left(\frac{L_c}{E_c} - 1\right)\right]^{1-\nu}}$ with respect to E_c/L_c . Thus, β_3 reflects the log-linear effect

of labor market tightness, as captured by the local employment rate (evaluated at parameter values such that all cities share the same employment rate), on the cost of filling a vacancy and is ultimately a function of parameters of the matching function, the cost of maintaining a vacancy, the discount rate, and the job destruction rate. We have written the equation in differences over time in order to eliminate time-invariant city-industry effects, and this is the way we will estimate it.¹⁷ In the data work, time periods will, for the most part, be ten years apart.

The important element in (5), relative to (3), is the presence of a negative feedback from the aggregate rate of employment to the rate of employment in one industry. This negative feedback, which reflects search externalities, may, at first pass, appear counterintuitive since one might expect that cross-good demand linkages would imply a positive feedback. However, for goods traded on a national market, the demand effects in our formulation should be captured by the industry-specific terms contained in β_{0i} , implying that the local aggregate employment rate captures the effect of search frictions.

B. The Intensive and Extensive Margins of Labor Demand

Equation (6) constitutes our baseline specification for labor demand at the industry city level as it relates hiring decisions to readily observable variables. This equation captures adjustments in employment demand along two distinct margins: the intensive margin (along which a given number of employers react to changes in labor costs) and an extensive margin (along which the number of active entrepreneurs is determined). For most policy questions the combined effect from both margins is the main object of interest, but it can nevertheless be useful to know the importance of each of these two margins separately. In this subsection we will derive simple analogs to (6) for each of the margins. We have previously noted that the intensive margin of employment is governed by the relation $E_{ic} = N_{ic}^a \cdot e_{ic}$, that is, employment along the intensive margin is the aggregation of individual-level employment decisions for all active firms. The extensive margin, on the other hand, reflects the decisions of potential entrepreneurs to become active, as captured by the relationship $N_{ic}^a = \gamma_{0i} G(f^*) L_c^{\gamma_1}$. Taking the same steps that allowed the derivation of (3) and then using (4) to replace the cost to a firm of hiring a worker with the combination of the wage received by the worker and a search cost, we can express the intensive margin of labor demand expression in log differences as

$$(7) \quad \Delta \ln E_{ic} \approx \beta_{0i}^I + \beta_1^I \Delta \ln w_{ic} + \beta_1^I \Phi \Delta \ln \frac{E_c}{L_c} + \beta_2^I \Delta \ln N_{ic}^a + \epsilon_{ic}^I,$$

¹⁷ Differencing also eliminates the fixed factor component of the error term since it does not vary over time by definition.

where $\beta_1^I = -\frac{(1 - \frac{\alpha_1 \alpha_3}{1 - \alpha_1 - \alpha_2} - \alpha_2 - \alpha_3)}{1 - \alpha_1 - \alpha_2 - \alpha_3}$, $\beta_2^I = 1 - \alpha_3$, $\epsilon_{ic} = \frac{1 - \frac{\alpha_3}{1 - \alpha_1 - \alpha_2}}{1 - \alpha_1 - \alpha_2 - \alpha_3} \times (\Delta \ln \theta_{ic})$, and β_{0i}^I again gathers industry-specific terms that are common across cities.

The extensive margin of adjustment, which can be referred to as the entrepreneurship equation, is accordingly represented by

$$(8) \quad \Delta \ln N_{ic}^a \approx \beta_{0i}^E + \beta_1^E \Delta \ln w_{ic} + \beta_1^E \Phi \ln \frac{E_c}{L_c} + \beta_2^E \Delta \ln L_c + \epsilon_{ic}^E,$$

where $\beta_1^E = -\frac{\phi \alpha_1}{1 - \alpha_1 - \alpha_2 + \alpha_3 \phi}$, $\beta_2^E = \frac{\gamma_1 (1 - \alpha_1 - \alpha_2)}{1 - \alpha_1 - \alpha_2 + \alpha_3 \phi}$, $\epsilon_{ic}^E = \frac{\phi}{1 - \alpha_1 - \alpha_2 + \alpha_3 \phi} \times (\Delta \ln \theta_{ic})$, and β_{0i}^E are industry-specific terms.

Relative to equation (6), the estimation of equations (7) and (8) can provide additional insight into the determination of employment demand. For example, by looking at the ratio of the coefficients on wages (β_1^I and β_1^E), we can get an estimate of the relative importance of the two margins in explaining how employment demand responds to changes in the cost of hiring workers. Note that each of the coefficients should be smaller in absolute value than the coefficient on wages estimated from (6). Furthermore, and possibly more importantly, by looking at the coefficients β_2^I and β_2^E one can evaluate the relevance of the mechanisms we have put forward on the role of entrepreneurs and of population in determining labor demand. In particular, our proposed mechanisms imply that entrepreneurship activity responds strongly to population (i.e., β_2^E should be large and significant) and that employment demand responds strongly to active entrepreneurship (i.e., β_2^I should also be large and significant). The most telling case would be if both these coefficients are close to 1, which would underline the key importance of entrepreneurship in the determination of employment demand as well implying that entrepreneurship tends to increase in proportion to population for fixed factor prices. Finally, by looking separately at the coefficients in these two equations, we can get a better sense of the role of fixed factors versus span of control problems in explaining a less than infinite elasticity of employment with respect to wages. Since the coefficient β_2^I should provide a direct estimate of the relevance of fixed factors, if this coefficient is close to 1, it would imply that a small estimate of the effect of wages on employment demand should be interpreted as mainly reflecting a combination of a span of control problem combined with the local supply of entrepreneurs being less than perfectly elastic.

C. Aggregating Industry-Level Labor Demand to City-Level Demand

Equation (6) expresses a labor demand curve at the industry-city level. Since policy analysis is more often focused on aggregate outcomes, as opposed to industry-specific outcomes, it is of interest to derive a city-level labor demand curve from (6). In the absence of search frictions, the aggregation of labor demand from the industry level to the city level can be done through simple adding up: there is no equilibrium interaction as wages are taken as given. However, in the presence of

search frictions, and even at given wages, the hiring decisions in one industry affect the hiring decisions in other industries through labor market congestion effects. We need to take this element into account in deriving a city-level labor demand. To this end, let us first define η_{ict} as the fraction of employment in industry i in city c (i.e., $\eta_{ict} = \frac{E_{ict}}{\sum_j E_{jct}}$). Now consider aggregating equation (6) using weights η_{ict} , and using the approximation $\sum_i \eta_{ict-1} \Delta \ln \frac{E_{ict}}{L_{ct}} \approx \Delta \ln \frac{E_{ct}}{L_{ct}}$, in order to get

$$(9) \quad \Delta \ln E_{ct} = \frac{1}{1 - \beta_3} \sum_i \eta_{ict-1} \cdot \Delta \beta_{0it} + \frac{\beta_1}{1 - \beta_3} \sum_i \eta_{ict-1} \cdot \Delta \ln w_{ict} + \frac{\beta_2 - \beta_3}{1 - \beta_3} \Delta \ln L_{ct} \\ + \sum_i \eta_{ict-1} \frac{\Delta \epsilon_{ict}}{1 - \beta_3}.$$

This equation expresses the change in the employment within a city as being negatively affected by the average wage change in the city ($\sum_i \eta_{ict-1} \cdot \Delta \ln w_{ict}$), and positively affected by the weighted sum of the β_{0it} . Notice that β_{0it} reflects a national-level effect associated with an industry. To express β_{0it} as a function of observables, we average (9) across cities (using the weights $1/C$, where C is the number of cities). This gives

$$\sum_c \frac{1}{C} \Delta \ln E_{ict} = \beta_{0it} + \beta_1 \sum_c \frac{1}{C} \Delta \ln w_{ict} + \beta_3 \cdot \sum_c \frac{1}{C} \Delta \ln E_{ct} + (\beta_2 - \beta_3) \cdot \sum_c \frac{1}{C} \Delta \ln L_{ct},$$

where we have used the assumption that $\sum_c (1/C) \Delta \epsilon_{ict} = 0$ since $\Delta \epsilon_{ict}$ reflects changes in comparative advantage.

The latter equation implies that β_{0it} can be written as

$$(10) \quad \beta_{0it} = \sum_c \frac{1}{C} \Delta \ln E_{ict} - \varphi_2 \sum_c \frac{1}{C} \Delta \ln w_{ict} + d_t,$$

where d_t is a year effect that is common across cities. The first two terms on the right side of the equation above can be approximated as the growth of employment in industry i at the national level, denoted $\Delta \ln E_{it}$, and the growth of wages in industry i at the national level, denoted $\Delta \ln w_{it}$. Thus, equation (10) indicates that the industry-specific intercept in (6) is approximately equal to the national-level growth in employment in the industry corrected for the average wage growth in the industry. Using (10), we can write the job creation curve at the city level as

$$(11) \quad \Delta \ln E_{ct} = d_t + \frac{1}{1 - \beta_3} \cdot \sum_i \eta_{ict-1} \cdot \Delta \ln E_{it} + \frac{\beta_1}{1 - \beta_3} \cdot \sum_i \eta_{ict-1} \Delta \ln \frac{w_{ict}}{w_{it}} \\ + \frac{\beta_2 - \beta_3}{1 - \beta_3} \Delta \ln L_{ct} + \tilde{\zeta}_{ct},$$

where $\tilde{\zeta}_{ct}$ is the error term given by $\sum_i \eta_{ict-1} \frac{\Delta \epsilon_{ict}}{1 - \beta_3}$.

Equation (11) now expresses cross-city differences in employment changes as a function of three main components. The first is a general growth effect captured by $\sum_i \eta_{ict-1} \cdot \Delta \ln E_{it}$, which reflects the notion that a city should have a better employment outcome if it is initially concentrated in industries which are growing at the national level. Second, we have a negative wage effect, which captures within-industry adjustments to a change in the cost of labor. This is given by the term $\sum_i \eta_{ict-1} \Delta \ln(w_{ict}/w_{it})$, which is large if a city experiences wage growth across industries that is higher on average than that experienced nationally. Since β_1 is negative, a high value of $\sum_i \eta_{ict-1} \Delta \ln(w_{ict}/w_{it})$ will result in lower employment outcomes in the city. The third term corresponds to a population growth effect. Finally, the error term reflects changes in the city's comparative advantage.

A comparison of equations (11) and (6) reveals an important difference in the wage coefficients in each. The coefficient on the city-industry specific wage change in equation (6) is the direct effect of a wage change on the employment rate in an industry-city cell holding the aggregate employment change in the city constant. This reflects the response of firms in an industry if that industry is too small to have a substantial effect on the overall equilibrium in the city. However, in general, we would expect that the immediate effect of a wage change in i , as captured in β_1 , would only be a first-round response. The decrease in employment in i would imply a less tight overall labor market in the city which would raise the value of a vacancy for entrepreneurs to an extent captured by β_3 . The resulting employment changes would then have further effects. The ultimate outcome of that process on total employment in the city is given by $\beta_1/(1 - \beta_3)$, which is the coefficient on the aggregated wage change in (11). Given that β_3 is predicted to be negative, the total impact of the wage change at the city level will be smaller than the direct, industry-specific effect, reflecting the self-correcting nature of the search externalities. Recall that $\beta_3 = \beta_1 \Phi$ and so, the ratio of the direct wage effect to that at the city level equals, $1/(1 - \beta_1 \Phi)$, making it a function of the determinants of Φ : the average tightness of the labor market, flow costs of hiring, etc. Thus, for example, larger flow costs of hiring imply that the congestion externality effect is larger and, as a result, the city-level elasticity is smaller relative to the direct effect.

II. Identification

In general, we would not expect ordinary least squares (OLS) to provide consistent estimates of the coefficients in equation (6), as the error term consists of changes in city-industry comparative advantage in production (the θ_{ic} terms). We expect changes in comparative advantage to be correlated with both changes in the wage in a given industry-city cell and with movements in the city-level employment rate. If worker migration decisions are based only on wages and employment rates then there may be no reason to expect a correlation between the change in the city size and the error term once we condition on wage and employment rate changes; that is, there would be no correlation if a productivity change is only of interest to workers to the extent it changes wages and the chance of getting a job. However, we allow for the possibility of a more direct connection, using instrumental variables related to each of the main right-hand-side variables.

The main pillar of our instrumental variable strategy will be to follow and extend ideas first presented in Bartik (1993) and used in many subsequent studies.¹⁸ The idea in Bartik (1993) is to work within a regional setting to construct instruments of the form $\sum_i \omega_{ict} \Delta Q_{it}$, where ω_{ict} are a set of weights specific to city c , and ΔQ_{it} is a change in the variable Q_i at the national level. In the specific case considered by Bartik, the weights are the beginning-of-period employment shares across industries within a city and ΔQ_i is the growth rate in employment at the national level in industry i between $t - 1$ and t . The result is a prediction of the end-of-period city employment rate based on the idea that if a particular industry grows or declines at the national level, the main effects from that change will be felt most in the cities that have the highest initial concentration in that industry. Note that this particular Bartik instrument is actually the first variable on the right side of our equation (11). Moreover, we can see from (11) that this instrument is potentially a good candidate for instrumenting the employment rate in equation (6) as, if β_2 is close to 1, then $\sum_i \eta_{ict-1} \cdot \Delta \ln E_{it}$ should be correlated with the change in the city-level log employment rate. We will call that instrument, Z_{1ct} .

Given our reliance on Bartik-type instruments, it is important to clarify the conditions under which they are valid.¹⁹ We will specify those conditions for Z_{1ct} first, then set them out in more general terms. Recall that the error term in (6) is given by $\Delta \epsilon_{ict}$ and corresponds to changes in local (industry-city level) productivity. It seems reasonable to be concerned that this error term is correlated with changes in the employment rate in the city. Now consider the potential correlation of this error term with Z_{1ct} . Since Z_{1ct} varies across cities, we are concerned with the cross-city correlation between it and the error term, which we can write as

$$\sum_c \frac{1}{C} \sum_i \eta_{ict-1} \Delta \ln E_{it} \Delta \epsilon_{jct} = \sum_i \Delta \ln E_{it} \sum_c \frac{1}{C} \eta_{ict-1} \Delta \epsilon_{jct}.$$

Taking the limit of the correlation as C goes to infinity implies that the instrument is asymptotically uncorrelated with the error term if

$$(12) \quad \text{plim}_{C \rightarrow \infty} \sum_c \frac{1}{C} \eta_{ict-1} \Delta \epsilon_{jct} = 0.$$

It is intuitive (and straightforward to show) that η_{ict-1} is a function of the values of the ϵ_{jct-1} s. Thus, this latter condition can be written in terms of the ϵ s, in which form it is equivalent to the following condition holding for all c and i :

$$(13) \quad \text{plim}_{C \rightarrow \infty} \sum_c \frac{1}{C} \epsilon_{ict-1} \Delta \epsilon_{jct} = 0,$$

where $\Delta \epsilon_{jct} = (\epsilon_{jct} - \epsilon_{jct-1})$. Thus, the validity of the instrument depends on a random walk-type assumption. This is clearly a stringent assumption, and we would like to be able to test it. This is possible if there is more than one instrument, allowing

¹⁸ See in, for example, Blanchard and Katz (1992).

¹⁹ This discussion builds on an earlier one in BGS. Goldsmith-Pinkham, Sorkin, and Swift (2017) also discuss the sources of variation in Bartik-type instruments and provide advice for practitioners.

for over-identifying tests of the underlying assumptions. This is precisely how we will proceed. We will take as a maintained assumption that the driving forces in the model, given by the set of ϵ s, satisfy the conditions for Bartik-type instruments to be potentially valid. We will then propose a set of such instruments and test the over-identifying restrictions to see if such an assumption is reasonable.

We view several of the features of this example as reflecting general characteristics of Bartik-type instruments: (i) the estimation is done in over-time differences; (ii) the error term often is a function of differences in productivities; and (iii) the weights (which we called ω s earlier) are plausibly functions of the lagged productivity levels.²⁰ From this, two lessons carry over to other implementations of Bartik-type instruments. First, validity of the instruments requires a random walk-type assumption, typically in terms of productivity processes. Second, the national-level change component of the Bartik instrument (the ΔQ) does not enter the asymptotic consistency condition. This is true because the validity of the instrument depends on cross-city correlations and the cross-city variation in the Bartik instruments comes from differences in the ω_{ic} vectors and not from ΔQ_i , which takes a common value across cities. This means that, asymptotically, there is no reason to worry about how city-level changes aggregate to a national value for Q . It is important, though, that this is an asymptotic statement that is based on an assumption that as the number of cities goes to infinity, industries are spread across many of them (i.e., there is no industry that operates only in one, or a handful of cities, as the number of cities gets large).

We now turn to discussing instruments of the Bartik form that are likely correlated with the change in wages. We have argued previously that labor supply shifters provide dubious instruments for the wage since they may be correlated with shifts in the supply of entrepreneurs. Hence, we need to turn to other forces that may drive wage changes. To this end, we draw on search and bargaining theory and exploit insights presented in BGS regarding the role of industrial composition in affecting workers' outside options and, through bargaining, wages. The idea in BGS is straightforward. Consider two identical workers who meet with potential employers in the same industry but in different cities. Upon meeting, the worker and employer can form a match and begin production or they can continue to search. With search frictions, a match will produce a bilateral monopoly, and workers and firms can bargain over the available match surplus to determine the wage paid. For the worker, the value of continuing to search serves as an outside option in the bargaining process. If there are frictions hindering perfect and costless mobility across cities, the value of continued search will depend on local labor market conditions. Within a local labor market, this value will depend, in part, on the expected quality of other potential matches and the expected duration of search. In particular, BGS show that when workers can potentially meet firms in any industry, the value of a worker's outside options will depend on the industrial composition of her city. Differences in local industrial composition will translate into differences in wages

²⁰For example, in what is commonly called the Ethnic Enclave instrument used in examining the impacts of immigration on a local economy, the concern is that immigrants move to the economy because of changes in productivity (captured, at least partially, in the error term). The ω s in that example correspond to the proportion of immigrants from some source country that were located in a given city in an earlier period. That distribution of immigrants is plausibly correlated with productivity in the city in the earlier period, and the identifying assumption is that those earlier productivity levels are uncorrelated with the changes in productivity in the sample period.

via bargaining, even if the tightness of the labor markets are the same, since higher outside options allow workers to capture more of the surplus. For example, workers in, say, the chemical industry should be able to bargain a higher wage if they live in a city with high-paying steel mills than if they live in a city where the steel mills are replaced with low-paying textile mills. We exploit this idea to justify two instrumental variables that will help to consistently estimate (6) and (11). The two instruments will be valid under the same assumption as we stipulated for Z_{1ct} .

Putting this idea a bit more formally, BGS show that in the context of a multiple sector search and bargaining model, industry-city wages, w_{ict} , will tend to be higher in cities where $\sum_j \eta_{jct} w_{jt}$ is higher (where w_{jt} represents wages in sector i at the national level and η_{jct} is the relative size of industry i at the city level). Note that $\sum_j \eta_{jct} w_{jt}$ proxies the outside options of workers and is higher in cities where workers are more likely to meet vacancies in high-wage industries.²¹ Notice that this is not a mechanical result since the ability of workers to switch industries implies that it would arise even if we just focused on other industries by dropping i when calculating the city average wage.

It is useful to decompose the movements in $\sum_j \eta_{jct} w_{jt}$ as follows:

$$(14) \quad \Delta \sum_j \eta_{jct} w_{jt} = \left(\sum_j \eta_{jct-1} (w_{jt} - w_{jt-1}) \right) + \left(\sum_j w_{jt-1} (\eta_{jct} - \eta_{jct-1}) \right).$$

Equation (14) indicates that for a worker in a particular city, outside options will increase over time if employment in that city is concentrated in industries where wages are increasing at the national level or if the worker is in a city where there is a shift in industrial composition toward relatively high-paying sectors. Importantly, BGS show that workers value each source of change in the value of outside options equally; a worker bargaining a wage in a given sector doesn't care whether her outside options change because of shifts in industrial structure or shifts in industry wages since all that matters is the expected wage in the city outside the current firm. In our empirical work, we use each component of the shifts in outside options to form the basis of an instrument for wages in (6) and (11).

We construct our first wage instrument, which we will call Z_{2ct} , based on the first term in (14),

$$Z_{2ct} = \sum_j \eta_{jct-1} (\ln w_{jt} - \ln w_{jt-1}).$$

BGS show that this instrument is a good predictor of wage growth at the industry-city level and give a formal justification for its relevance based on the wage bargaining story discussed previously. Importantly, Z_{2ct} varies across cities and obtains its variation entirely from the η_{jct-1} s (the initial period local industrial composition). As in our discussion of Z_{1ct} , the national-level wage changes are not relevant for our consistency considerations since they are common across cities. As such, Z_{2ct} will be

²¹ This formulation assumes that all searchers meet jobs in a sector in proportion to the size of the sector in the local economy. Tschopp (2015, 2017) extends this model by allowing workers in cells defined by a combination of occupation and industry to have differential mobility to other cells. She shows that this differential mobility is relevant for understanding wage setting but also that working with our simpler, industry-only specification generates similar overall effects of shifts in composition on wages. We work with the simpler specification here in order to highlight firm responses to wage changes rather than the specific mechanisms of wage setting.

uncorrelated with the error terms in (6) and (11) (and, hence, will be a valid instrument) under the assumption given in (13), that the comparative advantage terms, ϵ_{ict} , behave as random walks with changes independent of past levels.²²

The second instrument we propose for wages builds on the second term of (14), $\sum_j w_{jt-1} (\eta_{jct} - \eta_{jct-1})$. This term would not be an appropriate instrument since its dependence on the current industrial structure as captured by the η_{jct} s implies that it will not be orthogonal to the error terms in (6) or (11). Instead, consider the closely related variable given by

$$(15) \quad Z_{3ct} = \sum_j \ln w_{jt-1} \cdot (\hat{\eta}_{jct} - \eta_{jct-1}) = \sum_j \eta_{jct-1} \cdot (g_{jt}^* - 1) \cdot \ln w_{jt-1},$$

where $g_{jt}^* = \frac{1 + \Delta \ln E_{jt}}{\sum_k \eta_{kct-1} (1 + \Delta \ln E_{kt})}$. For the variable Z_{3ct} , we have replaced the current industrial composition term η_{jct} with its predicted value based on η_{jct-1} and the national-level trend in employment patterns.²³ As with Z_{1ct} and Z_{2ct} , the resulting variable's cross-city variation stems from the η_{jct-1} s and the same random walk assumption is needed for consistency. Furthermore, it should have predictive power for industry-city wage changes as it should capture the higher value of outside options for workers in a city where we predict that the industrial composition is tilting toward higher-paying jobs.

The availability of two instruments for wages raises the possibility of implementing an over-identification test. Variables Z_{2ct} and Z_{3ct} are both predicted to have an impact on city-industry wages through channels related to workers' outside options. But the channels that each exploits are quite different: one related to shifts in industrial structure and one to within-industry wage movements. As discussed above, theory predicts that each source of variation in outside options should have the same impact on wages since what matters for workers' bargaining positions is the change in the average wage in other industries, regardless of whether that change stems from changes in industrial composition or the industrial wage premia. Likewise, since what matters for employers is the bargained wage, variation in wages induced by either Z_{2ct} or Z_{3ct} should produce the same employment response. Since Z_{2ct} and Z_{3ct} rely on different forms of variation but are predicted to have the same employment impacts, this setup lends itself naturally to an over-identification test of the validity of our identification assumptions. Recall that both Z_{2ct} and Z_{3ct} are valid under the same random-walk assumption, that the η_{ict-1} s are uncorrelated with the $\Delta \epsilon_{ict}$ s in equations (6) and (11). If this assumption were violated, the offending correlations will be weighted differently by the two instruments (with changes in national-level industrial wages in Z_{2ct} and national-level employment changes in Z_{3ct}).

²²BGS presents a formal derivation of the form of the error term in the wage equation and proves that the conditions listed here imply that these instruments are valid.

²³To create the predicted share term, we first predict the level of employment for industry i in city c in period t as

$$\hat{E}_{ict} = E_{ict-1} \left(\frac{E_{it}}{E_{it-1}} \right).$$

Thus, we predict period t employment in industry i in city c using the employment in that industry-city cell in period $t - 1$ multiplied by the national-level growth rate for the industry. We then use these predicted values to construct predicted industry-specific employment shares, $\hat{\eta}_{ict} = \frac{E_{ict}}{\sum_i \hat{E}_{ict}}$, for the city in period t .

This would, in turn, imply that the two instruments should result in quite different estimated coefficients if the key correlations do not equal 0. Thus, we can test our identification assumption by testing that estimation of (6) and (11) using either Z_{2ct} or Z_{3ct} produces similar results. We view this test as quite strong because Z_{2ct} and Z_{3ct} work from quite different sources of variation; in fact, in our data their correlation is only 0.18 after removing year effects.

Recall that in equation (6) we have three explanatory variables for employment (besides the industries dummies). As we suspect all three of these variables to be potentially correlated with the error term, we need at least three instruments to estimate this equation. We have now proposed three instruments and, so, meet the order condition. But we are concerned that we may not meet the rank condition since the instruments presented so far are justified in relation to the wage and employment rate variables, not the third endogenous variable: population growth. For that reason, we now propose an instrument aimed at helping isolate admissible variation in the latter variable.

The population-related instrument is again of the Bartik form, and will be referred to as Z_{4ct} . The idea behind this instrument is to use historical patterns of interstate migration to predict inflows and outflows of people to a city.²⁴ For example, suppose a city has a large proportion of its population at the beginning of a period which is born out of state, young, female, and black. We infer that such a city is likely attractive to young, black females. Our proposed instrument is based on the prediction that such a city will grow if the out-of-state population of young, female, black people grows. To be more precise, Z_{4ct} is constructed as follows:

$$Z_{4ct} = \sum_j \omega_{jct-1} \cdot g_{jst},$$

where ω_{jct-1} is the fraction of the population in city c at time $t - 1$ that is both born out of state and is in demographic group j ; and g_{jst} is the growth between $t - 1$ and t of the out-of-state population in demographic group j . We segment the population into 40 demographic groups based on indicators for female, black, and age grouped into 5-year bins, using only those born in the United States.²⁵ Note that one of the sources of variation for this instrument is the aging of the baby boom, with this instrument predicting high population growth in cities where people of a given age group have tended to locate in the past as the baby boom moves through that age range.

A. Worker Heterogeneity

As we have emphasized, our aim in this paper is to provide an estimate of how employment decisions, on average, are affected by an across-the-board increase in the cost of labor. By its very nature, this question is about an aggregate labor market outcome. In the model developed so far workers are identical and so all parameters

²⁴This is very similar in nature to the ethnic enclave instrument that has been used in the literature estimating the effect of immigration on local labor markets starting with Card (2001). We use a slightly different form because we want to capture population effects relating to more than immigration since immigration may be a special type of shock (Peri and Sparber 2009).

²⁵The weights ω_{jct-1} in this case do not sum to 1.

are “aggregate” by definition. However, in our data, workers are heterogeneous in many dimensions including, among others, education and experience, and we need to address this heterogeneity in order to proceed appropriately. Depending on the assumptions that one makes, there are several ways to approach this issue.

The first approach, which we use for our main set of results, is to treat individuals as representing different bundles of efficiency units of work, where these bundles are treated as perfect substitutes in production. Therefore, in our baseline results we control for skill differences in wages via a rich regression adjustment and we correct for selection of workers across cities. This approach implicitly introduces an additional term in (6) which represents changes in average efficiency units per worker. In our baseline specification we treat this extra term as a part of the error structure, while in the robustness section we will show that our results are not sensitive to explicitly controlling for measures of efficiency units per capita at the local level. An alternative assumption is that labor markets are segregated along observable skill dimensions and that our model applies to homogeneous workers within these markets. Based on this, we also perform our analysis separately by education group as a specification check.

III. Data Description and Implementation Issues

Our main data come from the US decennial censuses for the years 1970 to 2000 and from the American Community Survey (ACS) for 2007, 2008, 2014, and 2015. We pool together the 2007 and 2008 ACS data for sample size reasons and also do the same thing for the 2014 and 2015 ACS years. For the 1970 Census data, we use both metro sample Forms 1 and 2 and adjust the weights for the fact that we combine two samples.²⁶ We focus on individuals residing in one of our 152 metropolitan areas at the time of the Census. Census definitions of metropolitan areas are not comparable over time. The definition of cities that we use in this paper attempts to maximize geographic consistency across Census years. Since most of our analysis takes place at the city-industry level, we also require a consistent definition of industry affiliation. Details on how we construct the industry and city definitions are left to online Appendix A.

As discussed earlier, our approach to dealing with worker heterogeneity is to control for observed characteristics in a regression context. Since most of our analysis takes place at the city-industry level, we use a common two-step procedure. Specifically, using a national sample of individuals, we run regressions separately by year of log weekly wages on a vector of individual characteristics and a full set of city-by-industry dummy variables.²⁷ We then take the estimated coefficients on the city-by-industry dummies as our measure of city-industry average wages, eliminating all cells with fewer than 20 observations. We adjust our standard errors to

²⁶ Our data were extracted from IPUMS: see Ruggles et al. (2015).

²⁷ We take a flexible approach to specifying the first-stage regression. We include indicators for education (four categories), a quadratic in experience, interactions of the experience and education variables, a gender dummy, black, Hispanic, and immigrant dummy variables, and the complete set of interactions of the gender, race, and immigrant dummies with all the education and experience variables.

account for wages being a generated regressor.²⁸ Our Z_{2ct} and Z_{3ct} instruments are constructed as functions of the national-industrial wage premia and the proportion of workers in each industry in a city. We obtain national-industrial premia by retaining the coefficients from regressions of the adjusted city-industry average wages on a set of industry dummies, estimated separately by year.

Our interpretation of the regression-adjusted wage measure is that it represents the wage paid to workers for a fixed set of skills. However, since we only observe the wage of a worker in city k if that worker chooses to live and work in k , self-selection of workers across cities may imply that average city wages are correlated with unobserved worker characteristics such as ability. In this case, our wage measure will not only represent the wage paid per efficiency unit but will also reflect (unobservable) skill differences of workers across cities. To address this potential concern, when we estimate our wage equations we control for worker self-selection across cities with a procedure developed and implemented by Dahl (2002) in a closely related context.

Dahl proposes a two-step procedure in which one first estimates various location choice probabilities for individuals, given their characteristics such as birth state. In the second step, flexible functions of the estimated probabilities are included in the wage equation to control for the nonrandom location choice of workers.²⁹ The actual procedure that we use is an extension of Dahl's approach to account for the fact we are concerned with cities rather than states, as in his paper, and that we also include individuals who are foreign born. When we estimate the wage equations, the selection correction terms enter significantly, which suggests that there are selection effects. Our results with or without the Dahl procedure are very similar. Nevertheless, all estimates presented below include the selection corrected wages.³⁰

The dependent variable in our analysis is the log change in industry-city level employment. We construct this variable by summing the number of individuals working in a particular industry. Our measure of L_c for a city is the city working-age population.³¹ For most of our estimates, we use decadal differences within industry-city cells for each pair of decades in our data (1980–1970, 1990–1980, 2000–1990) plus the 2007/8–2000 difference and the 2014/15–2007/8 difference, pooling these together into one large dataset and including period-specific industry dummies. In all the estimation results, we calculate standard errors allowing for clustering at the city level.

In our estimation of intensive and extensive margin responses, we also use data on net changes in establishments by industry and city cell from the Quarterly Census of Employment and Wages. We describe the details of matching that data to our city and industry definitions in online Appendix A.

²⁸ We derive the relevant correction for our asymptotic variance-covariance matrix in online Appendix D. Note that our instruments are also a function of the generated, adjusted wages but generated regressors as part of instruments do not require added adjustments to standard errors (Wooldridge 2002).

²⁹ Since the number of cities is large, adding the selection probability for each choice is not practical. Therefore, Dahl (2002) suggests an index sufficiency assumption that allows for the inclusion of a smaller number of selection terms, such as the first-best or observed choice and the retention probability. This is the approach that we follow.

³⁰ Details on our implementation of the Dahl's procedure are contained in online Appendix C. Results without the selection corrections are available upon request.

³¹ We have verified the robustness of our results to restricting the population to include only those individuals who report themselves as being in the labor force.

IV. Estimates of Labor Demand: Basic Results

A. Industry-City Level Results

In Table 1, we present estimates of our main equation of interest, (6). All the reported regressions include a full set of year-by-industry dummies. Column 1 reports OLS results. For the OLS results, both the coefficients on the wage and the city-level employment rate are positive and highly significant. This is the opposite of what our theory predicts. However, the employment equation derived from the model implies that OLS estimation of this equation should not provide consistent estimates. The fact that productivity shocks, $\Delta \epsilon_{ict}$, enter the employment equation's error terms, and that wages are likely positively related to productivity, explains why the OLS regression coefficient on wages is positive.

Column 2 contains results associated with estimating (6) using our full set of instruments, Z_{1ct} to Z_{4ct} . In the bottom rows of the table, we report the p -values associated with conventional F -tests of the joint null hypothesis that our four instruments have zero coefficients in the first-stage regressions associated with each of our endogenous variables. The p -values, in each case, are zero to two decimals, indicating that we do not have a weak instrument problem.

The first aspect to note about the IV results, relative to OLS, is that the IV coefficients on wages and the city-employment rate have the predicted negative sign. In particular, the coefficient on wages is estimated to be -1.03 , while the coefficient on the employment rate is estimated to be -2.09 , with both being statistically significantly different from 0 at the 5 percent level. The negative effect of the local employment rate on industry-level employment fits with the negative congestion effect suggested by search theory. For population changes, we find a strongly positive relationship, with a coefficient not significantly different from 1.³² Thus, holding wages constant, an increase in the labor force is associated with a close to proportional change in employment. Recall from Section II that a coefficient of population growth of 1 likely indicates that there are no important fixed factors at the industry-city level beyond that associated with a span of control problem. In the remaining columns of Table 1, we follow up on this result by imposing a coefficient of 1 on population growth. We implement this by using as our dependent variable the employment rate in an industry-city cell instead of the level of employment. Our results after imposing this restriction remain very similar. In particular, in column 6, which contains results using Z_{1ct} , Z_{2ct} , and Z_{3ct} as instruments, the estimates of the wage elasticity (β_1) and employment rate elasticity (β_3) are both statistically significant and take values of approximately -1 and -2 , respectively.

³²Manning and Amior (2015) document a strong persistence of employment rates by location in the United States over time and show that this is accompanied by (somewhat less strong) population inflows to high employment rate locations. They argue that these patterns fit with a model with gradual population adjustments to persistent employment shocks. Their estimation is based on a model with changes in population regressed on coincident changes in employment and the lagged employment rate and uses Bartik instruments. They obtain a coefficient on current employment that is less than 1. This seems to contradict our finding of a coefficient of 1 for the effect of population changes on employment rate changes. However, our result refers to a specification that controls for wage changes, which they do not include. Manning and Amior (2015) suggest that the differences in our results stems from the use of a different population variable but we use the same variable as they do.

TABLE 1—ESTIMATES OF LABOR DEMAND EQUATION (6)

	OLS (1)	IV (2)	OLS (3)	IV		
				(4)	(5)	(6)
$\Delta \log w_{ict}$	0.13 (0.013)	−1.03 (0.37)	0.12 (0.014)	−1.07 (0.29)	−1.04 (0.31)	−1.05 (0.27)
$\Delta \log \frac{E_{ct}}{L_{ct}}$	0.79 (0.054)	−2.09 (1.03)	0.79 (0.057)	−2.10 (1.06)	−2.22 (0.83)	−2.15 (0.81)
$\Delta \log L_{ct}$	0.89 (0.014)	0.99 (0.093)				
Observations	44,028	44,028	44,028	44,028	44,028	44,028
R^2						
Instruments		Z_1, Z_2, Z_3, Z_4		Z_1, Z_2	Z_1, Z_3	Z_1, Z_2, Z_3
F -stats						
$\Delta \log w_{ict}$		27.80		40.65	25.07	32.93
$\Delta \log \frac{E_{ct}}{L_{ct}}$		6.90		7.40	11.68	8.58
$\Delta \log L_{ct}$		19.70				
AP p -value:						
$\Delta \log w_{ict}$		0.00		0.00	0.00	0.00
$\Delta \log \frac{E_{ct}}{L_{ct}}$		0.00		0.00	0.00	0.00
$\Delta \log L_{ct}$		0.00				
Over-id. p -value		0.92				0.91

Notes: Standard errors, in parentheses, are clustered at the city-year level. All models estimated on our sample of US cities using Census and ACS data for 1970–2015 and include year fixed effects. The dependent variable is the decadal change in log industry-city employment (columns 1–2) or log industry-city employment rates (columns 3–6).

In columns 4 and 5, we repeat our estimation but include only Z_{2ct} or Z_{3ct} separately (along with Z_{1ct}) rather than both together as in column 6. Recall that these two instruments exploit very different data variation and so offer a good setup for exploring over-identification restrictions. In particular, if our identification assumptions are correct then we should get very similar results for the wage elasticity if we use either one of these instruments. This conjecture is confirmed as the wage elasticity is close to -1 using either set of instruments. The last row of Table 1 provides the p -value from a Hansen's J test of over-identification in which we test for differences between the estimated coefficients in column 4 and column 5. Not surprisingly, given the similarity of the estimated coefficients in the two columns, we do not reject the null hypothesis that the coefficients are the same across the different instrument sets. In fact, the p -value associated with the test is 0.91. We view the fact that our IV estimates are both changing the coefficients quite drastically compared to OLS results and are stable across instrument sets as strong support for our IV approach and the search theory underlying it. In BGS, we show the same sort of over-identifying result for wage equations and provide a more detailed interpretation. The other key prediction from the model is that an increase in labor market tightness in a city (as represented by the city-level employment rate) should negatively affect within industry-employment rates. Once we instrument, we do, in fact, find evidence of this negative effect. This is a striking result since it may have been

reasonable to expect a positive relationship between these variables. In our opinion, it is rather difficult to explain this later result without relying on search costs.

Separating the Margins of Adjustment.—As we discussed earlier, our estimates of the response of labor demand to changes in the cost of labor reflects both an intensive margin of adjustment (holding the number of entrepreneurs/firms fixed) and an extensive margin of adjustment (where the number of active entrepreneurs adjusts). To explore the relevance of each of these margins, we turn now to estimating equations (7) and (8). In order to estimate these equations, we need a measure or proxy for N_{ict}^a . Recall from the model that N_{ict}^a is what we call active entrepreneurs: that is, entrepreneurs with active firms with employees. It is difficult to find a direct measure of entrepreneurs of this type since surveys mainly ask about self-employment which could include what is essentially contract labor and people who incorporate for tax reasons rather than to create an employment generating firm. Measures of number of firms captures some of what we want but misses the individual entrepreneur element. In response, we use a combination of self-employment and enterprise data. The self-employment data come from our Census and ACS data and correspond to incorporated self-employed individuals in nonprofessional occupations. We choose incorporated self-employed as opposed to non-incorporated self-employed as these individuals are much more likely to employ others (Light and Munk 2015; Levine and Rubinstein 2013), and we use nonprofessionals in order to avoid doctors and accountants who are generally not job creators.³³ Our second measure is the number of establishments in the relevant industry \times city cell which we obtain from the Quarterly Census of Employment and Wages (QCEW). The QCEW reports establishment counts based on county and industry, subject to disclosure limitations to prevent the release of identifying information regarding single establishments. We believe that movements in a common factor in these two very different measures is a good proxy for movements in active entrepreneurship of the type our model emphasizes. Interestingly, there is a relatively strong element of commonality between the two measures as the cross-city correlation in their log changes is 0.4.

To capture the common movements of our two measures, we estimate a system of four equations: equation (7) estimated using the self-employment measure of N_{ict}^a on the right-hand side; equation (7) using the establishment measure of N_{ict}^a on the right-hand side; equation (8) using the self-employment measure as the dependent variable; and equation (8) using the enterprise measure as the dependent variable. We impose cross-equation restrictions such that the coefficients in each of the versions of equation (7) are restricted to be the same and the same for the versions of equation (8). This is effectively equivalent to using an average of the two measures in order to focus on their common movements. The other variables in (7) and (8) are the same as those used for the estimation of (6). Since the error terms in these two equations are changes in the productivity parameters, θ_{ict} , the instruments used for the estimation of (6) will remain valid under the same assumption.

³³ One drawback from using this measure is that the number of incorporated self-employed in many of the city-industry-decade cells is zero or very small. Since our specification is in logs, this forces us to drop close to one-third of the sample, and results in noisier estimates than in our base sample.

TABLE 2—THREE-STAGE ESTIMATION FOR SYSTEMS OF SIMULTANEOUS EQUATIONS, 1970–2015

	(1)	(2)	(3)	(4)
<i>Employment</i>				
$\Delta \ln w_{ict}$	0.24 (0.31)	0.11 (0.25)	0.088 (0.20)	0.085 (0.19)
$\Delta \log ER_{ct}$	−0.72 (3.53)	0.80 (1.42)	−0.39 (1.58)	−0.42 (1.37)
$\Delta \ln \text{Estabs}_{ict}$	1.07 (0.26)		0.99 (0.11)	1.00 (−)
$\Delta \ln \text{Self Emp}_{ict}$		0.84 (0.14)		
<i>Entrepreneurship</i>				
$\Delta \ln w_{ict}$	−1.04 (0.62)	−1.13 (0.66)	−0.86 (0.37)	−0.90 (0.27)
$\Delta \log ER_{ct}$	−1.47 (4.20)	−4.50 (2.59)	−2.26 (2.15)	−2.42 (1.57)
$\Delta \text{labor force}$	0.91 (0.27)	1.21 (0.26)	0.98 (0.15)	1.00 (−)
Observations	36,443	23,139	23,139	23,139
Instruments	Z_1, Z_2, Z_3, Z_4	Z_1, Z_2, Z_3, Z_4	Z_1, Z_2, Z_3, Z_4	Z_1, Z_2, Z_3, Z_4

Notes: Standard errors, in parentheses, are clustered at the city level. All models estimated on our sample of US cities using Census and ACS data for 1970–2015 and include a full set of industry \times year fixed effects. The dependent variable is the decadal change in log industry-city employment (panel 1) or decadal change in log industry-city entrepreneurs (panel 2). Column 1 measures entrepreneurship using establishment counts from the QCEW, column 2 measures entrepreneurship using incorporated self-employed from the Census/ACS, and columns 3 and 4 use both measures with cross-equation restrictions imposed as discussed in the text.

In Table 2 we report estimates of equations (7) and (8) from a three-stage least squares procedure implemented to account for the cross-equation correlation structure of the error terms.³⁴ The intensive margin (equation (7)) and extensive margin (equation (8)) estimates are reported in the upper and lower panels of the table, respectively. Our preferred specification is reported in column 3, where we use both measures of entrepreneurship, imposing cross-equation restrictions as described earlier. Columns 1 and 2 contain estimates using each of the entrepreneurship measures on its own as a point of comparison.

Focusing on the results in column 3, several interesting patterns emerge. First, the estimates of the effects of $\Delta \ln N^a$ in (7) and of $\Delta \ln L_c$ in (8) take values of 0.99 and 0.98, respectively. Both are strongly statistically significantly different from zero and not statistically significantly different from 1 at any conventional level of significance. This supports the ideas that entrepreneurship may be an important factor for understanding labor demand and that, even at fixed wages, entrepreneurship grows with population size. Together these two observations provide a coherent explanation for our finding of a coefficient close to 1 on $\Delta \ln L_c$ when estimating (6). The fact that the effect of $\Delta \ln N^a$ on labor demand is essentially 1 can be interpreted

³⁴The reported standard errors are clustered at the city level to account for potential serial correlation and are again adjusted to account for the wage being a generated regressor.

as supporting the notion that locally fixed factors likely play a limited role in the determination of employment demand since this coefficient should be an estimate of $1 - \alpha_3$. The observations that the IV estimates of β_2^I and β_2^E are not significantly different from 1 leads us to impose these restrictions in column 4, which yields estimates that are very similar to those in column 3 and a bit more precise.

It is interesting to note that the point estimates of the effects of wages in Table 2 are close to adding to what we observed for the combined effect when estimating (6). The estimated coefficients on wages also suggest that the extensive margin may be more important than the intensive margin, with the wage coefficient in the extensive margin results in column 4 being close to the combined effect of -1 while the intensive margin wage effect is small and not statistically significantly different from zero. It seems reasonable to expect that the intensive margin adjustments occur in the short run after a wage increase while the extensive margin adjustments are longer run in nature since they involve the opening and closing of firms. To the extent this is true, the fact that the estimated wage effect is significant in the entrepreneurship equation (the lower panel) but not in the intensive margin equation (the upper panel), may fit with findings that immediate employment reactions to minimum wage changes are much smaller than those in the longer run (Baker, Benjamin, and Stanger 1999 and Sorkin 2015).

When looking at the breakdown of the effect of the congestion externality on employment demand, as captured by the coefficient on the employment rate, it is hard to make any strong inferences given the imprecision of the estimates. Nonetheless, the point estimates also support the idea that the extensive margin of adjustment is likely more important than the intensive margin. Overall, we interpret Table 2 as providing more direct support for our view that: (i) firms face a span and control problem; (ii) that the incentive to create new firms is important for understanding the determination of labor demand; and (iii) that the creation of new firms, even at fixed wage and prices, likely grows with population size.

B. City-Level Results

We turn now to comparing the wage elasticity of labor demand at the industry-city versus city level. First, consider a wage increase in a particular industry, holding overall employment rates constant. If the industry in question is not large enough to have a significant impact on overall employment rates, the IV estimates in Table 1 imply a labor demand elasticity at the industry level of about -1 .

What about wage increases for a city as a whole? Since all industries will adjust employment downward in response to a general wage increase, there will be feedback effects on overall employment rates. Allowing these equilibrium effects to play out using our estimates of equation (6) implies a city-level labor demand elasticity of $\beta_1/(1 - \beta_3)$ or of -0.33 based on our estimates in column 2 of Table 1. In other words, since β_3 is predicted to be less than 0 in the presence of search frictions, overall wage increases in a locality have a built in dampening effect on employment responses because they simultaneously increase the availability of workers. In our model, this leads to reduced search costs for firms. Thus, our framework suggests that the city-level labor demand curve should be less elastic than the industry-level demand curve by a factor of $1 - \beta_3$.

TABLE 3—ESTIMATES OF THE AGGREGATE LABOR DEMAND EQUATION (11)

	OLS (1)	IV (2)	OLS (3)	IV		
				(4)	(5)	(6)
$\Delta \log w_{ict}$	0.11 (0.032)	−0.28 (0.091)	0.11 (0.032)	−0.29 (0.12)	−0.27 (0.10)	−0.28 (0.092)
$\Delta \log L_{ct}$	1.00 (0.010)	0.99 (0.031)				
Z_{1ct}	0.095 (0.046)	0.22 (0.064)	0.091 (0.039)	0.21 (0.047)	0.21 (0.038)	0.21 (0.038)
Observations	760	760	760	760	760	760
R^2						
Instrument set		Z_2, Z_3, Z_4		Z_2	Z_3	Z_2, Z_3
F -stats						
$\Delta \log w_{ict}$		46.99		78.87	51.38	69.80
$\Delta \log L_{ct}$		7.42				
AP p -value						
$\Delta \log w_{ict}$		0.00		0.00	0.00	0.00
$\Delta \log L_{ct}$		0.00				
Over-id. p -value		0.94				0.89

Notes: Standard errors in parentheses. All models estimated on our sample of US cities using Census and ACS data for 1970–2015 and include year fixed effects. The dependent variable is the decadal change in log industry-city employment (columns 1–2) or log industry-city employment rates (columns 3–6).

Recall that we can also obtain an estimate of the city-level demand elasticity through direct estimation of the city-level specification (11). Estimates of (11) are presented in Table 3, columns 1 and 2, with estimates where we use the employment rate as the dependent variable in columns 3–6. All estimations in the table contain a full set of year dummies, whose coefficients we suppress for brevity. Our IV estimates of the coefficient on log changes in average city wages, which represents an estimate of $\beta_1/(1 - \beta_3)$, range from -0.27 to -0.29 . The coefficients on labor force growth are very close to 1, regardless of whether we estimate by OLS or IV. In the last three columns of the table, the wage elasticity obtained using Z_{2ct} and Z_{3ct} are again nearly identical to each other, and the over-identification test again fails to reject the null hypothesis associated with these being valid instruments (in this case, the test has an associated p -value of 0.89). Thus, in this city-level specification, the results continue to support our proposed framework for studying labor demand. It is important to emphasize that the estimates of the city-level demand elasticities using the aggregated data are almost identical to what we just calculated using the estimated coefficients from the industry-level specification (6). Since estimation of (6) and (11) use very different levels of aggregation, and since there is no mechanical reason the two specifications should provide the same results for $\beta_1/(1 - \beta_3)$, we view the similarity of the estimates of the city-level elasticity obtained from the two different approaches as evidence supporting our framework.³⁵ Finally, note that in Table 3, Z_{1ct} has a positive and strongly statistically significant direct effect on the

³⁵ Note that the OLS estimate of $\beta_1/(1 - \beta_3)$ obtained from the city-level estimation (0.11) is not close to that obtained from the OLS estimates of the industry city equation (which equals 0.62). This supports the claim that there is no obvious mechanical link forcing a similar result from the two estimates.

city-level employment rate, supporting the idea that it is a good instrument for that employment rate in the disaggregated equation estimation presented in Table 1.

C. Breakdown between Traded and Non-Traded Goods

In our model and interpretation of the data, we have assumed that all goods are traded across cities. This assumption allows us to treat the price of goods as being common across cities and, therefore, to fully capture their effects through time-varying industry effects. If there are goods produced that are not tradable across cities, it will create a city-specific component in prices that will appear in the error term of our labor demand regressions. A simple way to get around this problem is to focus only on labor demand in tradable goods sectors. To this end, we define tradable and non-tradable sectors using an approach from Jensen and Kletzer (2006). They argue that the share of employment in tradable goods should vary widely across regional entities (cities in our case) since different cities will concentrate in producing different goods which they can then trade. For non-tradable goods, on the other hand, assuming that preferences are the same across cities, one should observe similar proportions of workers in their production across cities. We therefore rank industries by the variance of their employment shares across cities in the 1970 Census and label the industries in the top, middle, and bottom third as high-, medium-, and low-trade industries.³⁶

In Table 4, we present estimates of equation (6) carried out separately for the low-, medium-, and high-trade industries with the coefficient on population constrained to be 1. Estimates derived without constraining that coefficient are given in online Appendix F, and they show that the population coefficient is not significantly different from 1 for any of the groups. The odd-numbered columns of Table 4 report OLS estimates, while the even-numbered columns report IV estimates. The IV estimates mirror our overall results in showing impact (β_1) estimates that are close to -1 and substantial congestion externality effects in all three groups. But they are not identical. The implied long-run wage cost elasticities are -0.32 , -0.27 , and -0.42 for the low-, medium-, and high-trade groups, respectively. This fits with intuition that one would expect a higher elasticity in high-trade industries since, effectively, there are more substitutes available for the final goods in those industries.

To push potential industry differences further, in Table 5 we report estimates of the wage elasticity of labor demand for seven common industry groupings. In the first column of this table we report OLS estimates of this elasticity and in the second column we report IV estimates. For these estimates we have constrained the coefficient of population growth to equal 1.³⁷ The estimated wage elasticity is negative in all industries and is only not statistically significant in agriculture, mining, and construction. For the other six industry groupings, wage elasticities range between

³⁶Note that our interest in this exercise is not in the potential effects of trade shocks in shifting labor demand, something that has been examined using similar regional variation to ours by Autor, Dorn, and Hanson (2013). Rather, we are interested in focusing on traded goods in order to cut out local, non-traded good price changes. In our context, the kinds of trade-induced demand shifts examined in Autor, Dorn, and Hanson (2013) will be captured partly in the general demand shifts reflected in our Bartik instrument for ER_i and partly in the instruments for the average wage in a city that are based on predicted shifts in the industrial composition.

³⁷We have omitted the estimates on the city-level employment rate to save space.

TABLE 4—ESTIMATES OF LABOR DEMAND EQUATION (6) BY TRADE GROUPS

	Low-trade		Medium-trade		High-trade	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta \log w_{ict}$	0.13 (0.026)	-1.24 (0.56)	0.090 (0.020)	-1.14 (0.37)	0.15 (0.025)	-0.93 (0.20)
$\Delta \log \frac{E_{ct}}{L_{ct}}$	0.66 (0.13)	-2.84 (2.58)	0.83 (0.078)	-3.23 (1.18)	0.81 (0.088)	-1.22 (0.57)
Observations	6,824	6,824	18,469	18,469	18,735	18,735
R^2						
Instrument set		Z_1, Z_2, Z_3, Z_4		Z_1, Z_2, Z_3, Z_4		Z_1, Z_2, Z_3, Z_4
F-stats						
$\Delta \log w_{ict}$		9.63		27.50		46.88
$\Delta \log \frac{E_{ct}}{L_{ct}}$		2.28		7.52		12.05
AP p -value						
$\Delta \log w_{ict}$		0.00		0.00		0.00
$\Delta \log \frac{E_{ct}}{L_{ct}}$		0.04		0.00		0.00
Over-id. p -value		0.77		0.99		0.86

Notes: Standard errors, in parentheses, are clustered at the city-year level. All models estimated on our sample of 152 US cities using Census and ACS data for 1970–2007. The dependent variable is the decadal change in log industry-city employment rates.

TABLE 5—WAGE ELASTICITY ESTIMATES BY INDUSTRY AGGREGATES

	OLS (1)	IV (2)
Agriculture, mining, construction	0.20 (0.069)	-0.51 (0.38)
Manufacturing	0.21 (0.036)	-1.51 (0.48)
Transport, communications, utilities	0.036 (0.041)	-1.18 (0.54)
Retail, wholesale	0.10 (0.021)	-0.78 (0.25)
F.I.R.E.	0.088 (0.030)	-0.91 (0.29)
Personal, entertainment	0.046 (0.032)	-0.97 (0.26)
Professional	0.15 (0.036)	-0.97 (0.36)
Observations	41,932	41,932
R^2		
Instruments		Z_1, Z_2, Z_3
Average	0.14	-1.05

Notes: Standard errors, in parentheses, are clustered at the city-year level. All models estimated on a sample of 152 US cities using Census and ACS data for 1970–2007. The dependent variable is the decadal change in regression adjusted city-industry wages.

-0.78 and -1.51 . The average of the IV estimates using the industry shares as weights, reported in the last row of the table, is -1.05 . Hence, it appears reasonable to conclude that the wage elasticity of labor demand at the industry-city level is close to -1 .

V. Robustness

We have explored the robustness of our results along various dimensions. In online Appendix F, we present results separately for workers with high school or less education and those with some post-secondary or more. The results for the high-school-educated indicate a higher industry city impact of a wage cost increase than for the post-secondary-educated but both imply similar longer-run elasticities of about -0.30 once congestion externalities are taken into account. In online Appendix F, we also present an exercise in which we include lagged wage effects to allow for the possibility of dynamic effects that extend over periods of more than ten years. In particular, we reestimated our labor demand equation allowing for the initial level of wages to affect the change in employment, instrumenting for the initial wage with the wage level ten years prior. While the first stage is strong, we do not find substantial lagged wage effects and our main results are unchanged. Although this does not imply that other types of dynamic effects are not present, it does provide some support for the claim that our rather static specification of labor demand may be appropriate for studying changes in employment over decades.

VI. Summary and Interpretation of Empirical Results

From our estimation of equation (6) using data over four decades, we have found strong support for the following three patterns. First, we found a significant and robust negative wage elasticity of labor demand. This wage elasticity is estimated to be close to -1 at the industry-city level and -0.3 at the city level. Second, holding wages constant, an increase in the size of the labor force is associated with an increase in employment in a proportion close to one-to-one. Finally, tighter labor markets at the city level reduce industry-level employment within the city.

How should we interpret these results? The finding we believe to be most interesting is the joint observation of a wage elasticity of labor demand far from infinity combined with an estimated elasticity of the labor demand to population close to 1. The model presented in Section II suggests that one should infer from the latter observation on the effects of population that fixed inputs such as, for example, land, are unlikely to be placing important constraints on employment at the local level. This estimated population effect therefore also implies that the non-infinite wage elasticity of labor demand we estimate should not be interpreted as reflecting decreasing returns to scale due to some fixed physical factor. Recall from equation (3) that, in the absence of a fixed physical factor, the wage elasticity of labor demand should be equal to infinity if there is either an infinite supply of potential entrepreneurs or if there is no span of control problem.³⁸ Hence, observing a far from infinite

³⁸ As noted previously, our approach does not allow us to differentiate evidence of decreasing returns at the firm level between a span of control problem or a limited demand for differentiated goods.

wage elasticity of labor demand combined with a proportional effect of population on employment implies that there is a limited supply of entrepreneurs willing to open shop in response to profit opportunities and that those entrepreneurs face span of control problems within their firms. We emphasize this finding because it is rather common in the macro-labor literature to assume that the supply of entrepreneurs is infinitely elastic with respect to any profit opportunities. Our extensive margin estimates, instead, imply an elasticity of the supply of active entrepreneurs (i.e., entrepreneurs who open firms that take part in job creation) with respect to the wage they pay of -1 . While this elasticity is large, it indicates that assumptions of an infinite supply elasticity are unwarranted even when looking over rather long time spans.

It is interesting to reconsider the wide range of available estimates of the elasticity of labor demand in light of these results. At one extreme, studies examining local labor market effects of migration-related supply shocks tend to find large increases in the number of workers employed in the receiving labor market but little change in wages. This is what David Card found in his famous study of the 7 percent increase in the population of Miami generated by the Mariel Boatlift. In a standard neoclassical framework, this could be interpreted as implying a nearly perfectly elastic labor demand curve. However, we argue that the population inflow would likely bring with it more entrepreneurs and that this, alone, would imply an increase in employment.³⁹ Importantly, in our very general specification the resulting wage and employment changes cannot be used to identify the effects of a wage change on employment at the level of the local labor market. Instead, one would need to focus directly on mechanisms for generating reliable variation in wage costs. This is the goal of the minimum wage literature, but, there, one might be worried that the resulting estimates are specific to the low wage labor market. We, instead, make use of insights from the search and bargaining literature to obtain identifying variation based on wage spillovers from changes in the industrial composition of a city. The resulting estimates indicate that the city-level labor demand curve is much less than perfectly elastic.

The second insight we believe should be taken from our estimates is the relevance of search frictions. Allowing for such frictions in the estimation of labor demand curves has certainly not been the norm. However, our results suggest they are important. In particular, we saw from our estimates of industry-city level labor demand curves that, holding wages constant, employment at the industry level decreased when employment at the city level increased. Viewed through the lens of our framework, this pattern also implies that the wage elasticity of labor demand at the city level should be smaller than that at the industry level, which is precisely what we found. While there may exist other explanations for such a pattern, search frictions offer a simple rationalization of the observed effects. In summary, our framework explains the rather small wage elasticity of labor demand that we estimate at the city level as reflecting a combination of three factors: decreasing returns to scale at the firm level, limited supply of entrepreneurs, and search externalities.

³⁹The idea that immigrant inflows could bring with them flows of complementary factors of production goes back at least to Mutti and Gerking (1983). Our results in Beaudry, Green, and Sand (2014) indicate that increases in a city's population do not alter the wage rate over the medium term. Combined with the employment rate results here, this indicates that a search model with proportional entry of entrepreneurs can rationalize the findings in Card (1990a).

Our quasi-structural approach does not allow us to identify the exact extent of the decreasing returns to scale but does provide us with direct estimates of the supply elasticity of active entrepreneurs (-1), and of the extent of search externality effects (which are given by our estimates of β_3 and are large). Recall that we earlier wrote β_3 as a function of underlying structural parameters. Our derived empirical specification indicates that it is that function that captures the actual congestion effect and, so, we view it as more interesting than the underlying parameters.

VII. Application: The Employment Effects of a \$15 Minimum Wage

The movements in several large US cities to introduce a \$15 per hour minimum wage has renewed interest in the merits and costs of higher minimum wages. Leading examples of this movement include Seattle, which decided in 2014 to increase its minimum wage from the prevailing level of \$9.34 to \$15 beginning in 2017, and San Francisco, where residents voted to increase their minimum wage from \$10.74 to \$15 by 2018. In 2015, Los Angeles decided to follow suit, raising its minimum wage from \$9 to \$15 by 2020. The primary goal of these changes is to help lift working families out of poverty. Balanced against any such benefits are potential costs, particularly in terms of reductions in employment. In this section, we use our estimates of city-level wage elasticities to assess the likely employment effects of the moves to a \$15 minimum wage in Seattle, San Francisco, and Los Angeles. We take account of the fact that the increase will vary by firm size and be staged differently over time in the three cities. In addition, the fact that these cities have different wage distributions and industrial compositions will lead to different employment implications in each case.

Before describing our estimated impacts in detail, it is worth asking why one would not simply use the large existing set of estimates of the elasticity of employment with respect to the minimum wage. The first weakness we see with those estimates relates to the mechanics of how the effects in that literature are generally estimated. In most papers, researchers regress changes in employment rates on changes in minimum wages with varying sets of controls for local labor market conditions, etc. This means that the estimated employment effects should be seen as reduced-form entities which aggregate across two underlying steps in the way minimum wages affect employment. The first step involves calculating how a minimum wage change alters the cost of labor to firms while the second translates that cost change into employment effects at both the firm intensive and extensive margins. Given this, standard reduced-form estimates will obviously vary depending on where the minimum wage change fits into the distribution of wages. For example, if the minimum wage remained below the lowest wage paid in the market even after being raised then a standard estimate would, by construction, equal 0 regardless of how responsive is employment to changes in labor costs. This does not raise problems if the minimum wage change being considered is small and fits in the range of minimum wage changes that have been used to estimate the employment response elasticity. But the moves to \$15 minimum wages represent large changes that are outside the range of variation used in obtaining existing minimum wage effect estimates. Seattle's minimum wage change, for example, involves moving the minimum wage from the eighth to the twenty-eighth percentile of the overall

wage distribution, which sits much higher up the wage distribution than recent minimum wages in any jurisdiction in the United States.⁴⁰ Accordingly, it is not clear to us how to use estimates from the existing minimum wage literature to extrapolate employment effects from a change which affects a large segment of the wage distribution.⁴¹ In contrast, our estimates are aimed at directly capturing how a change in labor costs affects employment, and therefore can be used to evaluate the impact of small or large changes in the minimum wage as long as the induced change in the average wage is within our sample variation.

A second weakness of the existing minimum wage literature is its focus on specific submarkets. In particular, the vast majority of existing estimates are for teenagers and/or the restaurant sector.⁴² But our results indicate an important role for congestion externalities in the impact of changes in labor costs on employment in an economy. Thus, extrapolating from estimates for teenagers or one sector to impacts for the economy as a whole is problematic. In contrast, our predictions take direct account of congestion externalities. Once again, if a proposed change in the minimum wage is very targeted or small, such general equilibrium effects are likely small and therefore may be reasonably ignored. However, in the case at hand, the proposed minimum wage changes will have broad-based and substantial impacts, and therefore it is likely important to take into account any general equilibrium effects.⁴³

Finally, given that we work with differences across Censuses, our estimates capture medium- to long-run effects of changes in wage costs on employment rates. While many studies in the minimum wage literature try to capture long-run effects, the time frame examined varies greatly. Many natural experiment type papers focus on rather short-term variation of one or two years. Studies that exploit time series variation in minimum wages often include a lagged term to capture dynamics. However, it is well known that such an approach to capturing dynamics has pitfalls.

⁴⁰The \$15 minimum wage, even taking into account the fact that inflation will eat into its impact by the time it is actually implemented in most cities, lies above the range of much of the variation in minimum wages that has been used in existing research. Evans and Zipperer (2014) plot the ratio of the minimum wage to the median hourly wage for paid employees working at least 35 hours per week by state from 1980 through 2014. In the most recent 20 years, the weighted average of that ratio across states has stayed persistently near 40 percent and there are only one or two years in which any state reaches 50 percent. In comparison, in Reich et al.'s (2014) evaluation of the LA minimum wage (which is scheduled to reach \$15.25 in 2019), the minimum wage will reach 66.5 percent of the full-time median hourly wage by 2019 under their medium wage growth scenario. Similarly, using data from the 2013 Current Population Survey (CPS) outgoing rotation sample for Seattle, we find that the \$15 minimum wage in that city will be at 63 percent of the median when it is first implemented.

⁴¹This criticism does not apply to structural model applications such as those in Flinn (2006) and Ahn, Arcidiacono, and Wessels (2011) but those papers face the problem of focusing on a narrow subset of workers and not taking into account the type of general equilibrium interactions we associated with congestion externalities.

⁴²See Belman and Wolfson (2014) for a recent, thorough review of the minimum wage literature. The existing literature focuses on groups that are perceived to be most vulnerable to small shifts in the existing minimum wage (teenagers, restaurant workers, single mothers, immigrants, etc.) while a move to \$15 will be relevant for a much wider range of workers. Very few, if any recent studies present estimates for workers over age 30, perhaps because when minimum wage effects are estimated for older workers they are invariably economically insubstantial (Brochu and Green 2013). But it is unclear whether they are insubstantial because minimum wages truly don't affect these workers (i.e., they face very inelastic demand curves) or because minimum wages in the ranges witnessed in previous data have no "bite" on their wage distribution.

⁴³In online Appendix Table B, we present proportions by various characteristics for Seattle and Los Angeles workers with wages near the current minimum wage and near \$15. Those proportions indicate the broader effects of the higher minimum. For example, in Seattle, 17 percent of those earning within \$1 of the current minimum wage are teenagers and 30 percent work in the restaurant industry. In contrast, only 1 percent of workers with wages between \$14 and \$15 are teenagers and 10 percent work in the restaurant industry.

Sorkin (2015) points out that if the nominal minimum wage is increased but then left unchanged for several years afterward then the real minimum wage change is, in essence, temporary and firms will not change their capital-labor ratios much. In that situation, estimates of long-run effects that are obtained as the coefficient on lags of the minimum wage in reduced-form specifications will be negligible. But if (as is the case in many of the \$15 minimum wage policies) the minimum wage is pegged to inflation after being set at its new level then there will be more adjustment and, eventually, a greater loss of employment. Given this, it is relevant that our results are based on decade-level variation in wages and employment. Thus, our estimates allow time for adjustments at both the intensive (hiring by existing firms) and extensive (firm entry and exit) margins. This fits with the putty-clay model in Sorkin (2015).⁴⁴

A. Implementation

In our baseline model, wages at the industry-city level are assumed to be set in efficiency units and workers only differ in terms of the number of efficiency units they represent. However, to explore the effects of minimum wages, it appears more reasonable to think of workers as differing more than simply in terms of efficiency units. In particular we will think of workers as being of one particular skill type among Q possibilities. There is a market for each of these types of workers and entrepreneurs can hire them to produce goods, which are skill- and industry-specific, that are sold on the national market. In this extended setup, the relevant unit for the determination of employment and wages is at the skill-industry-city level while the price of the good produced is determined at the national level. Under these slightly modified assumptions, one can use the same steps as presented previously to derive the demand for a skill of type q in industry i in city c , which will be given by ⁴⁵

$$(16) \quad \Delta \ln \frac{E_{qic}}{L_{qc}} = \beta_{0i} + \beta_{q1} \Delta \ln w_{qic} + \beta_{q3} \Delta \ln \frac{E_q}{L_{qc}} + \epsilon_{qic}.$$

Equation (16) is well suited to examine the effects of an increase in minimum wage, as we can associate each different skill group with its wage payment. Accordingly, the effect of a minimum wage increase can then be calculated as being the employment effect of increasing the wages of the skill groups which previously had wages below the new minimum wage. Recall that this effect would be capturing adjustments at both the intensive and extensive margin. Such a calculation would be quite straightforward to carry out if we had estimates of β_{q1} and β_{q3} for each skill group. However, we do not have such a detailed set of estimates. To get a handle on how much these β s may vary across skill groups, we can first consider the case of only two skills: high-school-educated workers and college-educated workers. In online Appendix Table E10, we report estimates of equation (16) allowing for

⁴⁴Meer and West (2016) also point to the usefulness of examining long differences in order to fully capture firm adjustments.

⁴⁵In this specification, we are assuming that each of these markets exhibit constant returns to market size.

these two skill groups. The key outcome of this exercise is that the estimated complete wage elasticity, $\beta_{1q}/(1 - \beta_{3q})$, is nearly identical for the two education groups. Estimates using alternative groupings (for example, defined by gender and education) yield similar invariance results. From this we conclude that it is reasonable to make our calculations under the assumption that β_{1q} and β_{3q} do not vary with q and we base our actual calculations on that assumption.⁴⁶ However, to give a slightly richer structure to these calculations, we will allow β_1 to differ by broad industry groupings as reflected in Table 5. Although we will be using the same estimates of the β s across skill groups, our calculations will nevertheless imply that different skill groups (defined by their initial wages) will be more affected by the minimum wage than others. Accordingly, we will complement our estimates of city wide effect with results on the distributional effect for workers with different initial period wages.

Our evaluation of the impact of the different proposed increases in the minimum wage on employment uses individual data from the Current Population Survey's Outgoing Rotation Groups and proceeds in three steps.⁴⁷ First, for each city, industry, and skill we calculate how the minimum wage increase is likely to increase the wage. Second, we use our estimate(s) of β_1 at the industry level to calculate how the resulting increase in labor cost would reduce employment for each skill group in the absence of any feedback effects due to changes in overall market tightness (where the skill groups are defined by the initial wage). Third, we use our estimate of β_3 to adjust our skill level estimates to incorporate the general equilibrium effects induced by changes in overall market tightness. Combining the three steps, aggregating over all the skill groups, yields our estimate of the impact of the increase in the minimum wage on the employment rate in a city. Given our estimates of β , it is very simple to perform steps two and three. The main issues arise in step one as this is where differences between cities come into play.

In order to calculate the effect of a minimum wage rise on the cost of labor in a city, we construct a counterfactual in which workers with wages below the new minimum have their wages replaced with that new minimum. For other, above minimum earners, we do not alter their wages. That is, we assume that the rise in the minimum wage does not have any spillover effects on the wages of workers paid above the new minimum. This fits with many previous estimates of the impact of minimum wages on the wage distribution (e.g., Dickens and Manning 2004). In online Appendix B, we discuss results when using the wage setting model in BGS, in which wage spillovers are possible. Our approach generates wage changes that differ by initial wage level. For example, if a worker was previously paid \$12/hour, then a minimum wage change to \$15 induces a wage increase of 25 percent, while a worker previously paid \$14.50/hour only receives an increase of 3.4 percent ($0.5/14.5$). As a result, our estimates will differ with the initial wage distribution.

The only minor implementation complication is that the targeted minimum wages are being phased in gradually over time and in different manners across cities. We present details of how we calculate the wage effects of the planned minimum wage

⁴⁶ While this assumption may appear somewhat extreme, a similar assumption is generally being made when using natural experiment type estimates of the effects of the minimum wage, which are generally based on small changes in certain particular markets, to extrapolate large changes affecting many markets. The advantage of our approach is that our estimates of β_1 and β_3 are based on a large segment of the population.

⁴⁷ Details on data processing can be found in online Appendix A.A2.

TABLE 6—MINIMUM WAGE ROLL-OUT

	Large firms		Small firms	
	Nominal	2014 dollars	Nominal	2014 dollars
<i>Seattle</i>				
2014	9.32	9.32	9.32	9.32
2015	11.00	10.78	11.00	10.78
2016	13.00	12.50	12.00	11.53
2017	15.00	14.13	13.00	12.25
2018	15.30	14.13	14.00	12.93
2019	15.61	14.13	15.00	13.59
<i>Los Angeles</i>				
2014	9.00	9.00	9.00	9.00
2016	10.50	10.09	10.00	9.61
2017	12.00	11.31	10.50	9.89
2018	13.25	12.24	12.00	11.09
2019	14.25	12.91	13.25	12.00
2020	15.00	13.32	14.25	12.65
2021	15.00	13.06	15.00	13.06
<i>San Francisco</i>				
	All firms			
	Nominal	2014 dollars		
2014		10.74	10.74	
2015		11.05	10.83	
2016		12.25	11.77	
2017		13.00	12.25	
2018		14.00	12.93	
2019		15.00	13.59	

increases in Seattle, San Francisco, and Los Angeles in online Appendix B. In all cases, we first inflate the wage distribution from our most recent data (2014) to the year in which the \$15 minimum takes effect using a 2 percent inflation rate.⁴⁸ Since the year in which the new minimum takes effect varies between 2017 and 2021, this immediately implies that the employment effects are smaller than would be the case if the minimum wages were moved to \$15 in the current distribution. In the case of Seattle and Los Angeles, we also take into account the fact the minimum wage differentially affects workers at large and small firms.

Estimated Effects of the Minimum Wage.—In Table 6, we present the nominal and real minimum wage schedules for Seattle, Los Angeles, and San Francisco. For both Seattle and Los Angeles, the phasing in of the \$15 minimum wage differs for small versus large firms, with large firms being over 500 workers in Seattle and over 25 workers in Los Angeles. The \$15 minimum does not cover all workers until 2019 in Seattle and 2021 in Los Angeles. In Table 7, we show how the minimum wage phase-in affects the average log wage and the proportion of workers with wages at the minimum wage in Seattle under our assumption of no wage spillover effects.

⁴⁸ We chose 2 percent because the US Federal Reserve has committed to a 2 percent inflation target. Given the long periods over which the new minimum wages are phased in, our results are sensitive to the inflation rate and we provide some related robustness checks in online Appendix B.

TABLE 7—WAGE IMPACT FROM MINIMUM WAGE CHANGES: SEATTLE

	2015 (1)	2016 (2)	2017 (3)	2018 (4)	2019 (5)
1. Initial	3.04 (0.016)	3.06 (0.016)	3.08 (0.015)	3.11 (0.014)	3.12 (0.014)
2. Direct impact	0.021 (0.0013)	0.025 (0.0014)	0.028 (0.0013)	0.011 (0.00070)	0.011 (0.00068)
3. End-of-year	3.06 (0.016)	3.08 (0.015)	3.11 (0.014)	3.12 (0.014)	3.13 (0.014)
4. Fraction impacted	0.16 (0.0097)	0.21 (0.011)	0.27 (0.011)	0.28 (0.012)	0.30 (0.012)
Total wage change					0.098
Standard error					0.005

Notes: Wage Impacts calculated as follows: row 1 gives the average log wage at the beginning of the period. Row 2 gives the impact on the average log wage caused by the period's minimum wage increase. Row 3 gives the end-of-year average log wage (row 1 + row 2). Row 4 gives the cumulative fraction of workers impacted by the roll-out of the minimum wage policy. All wage figures are in 2014 dollars.

The annual changes in the average log wage are between 1 and 3 percent, with the ultimate impact being a 10 percent increase in the average wage.

In Table 8, we report estimates of the employment-rate impact of the minimum-wage-induced changes in the wage distribution for Seattle for each year until full phase-in.⁴⁹ The top-left entry in the table shows the initial employment rate in the city before the first step in the minimum wage schedule that takes place in 2015. The second row shows that the direct effect of the change from a \$9.32 to a \$10.78 minimum wage without accounting for congestion externalities is a decline in the employment rate by 1.4 percentage points. Row 3 shows that once congestion externalities are taken into account, the net impact of this initial wage change on the city employment rate is a decline of 0.4 percentage points. While this initial \$1.46 increase in the minimum wage is large by historical standards, it is close to variation that has been used in the existing literature. The small size of the resulting employment rate effect can then be seen as fitting with other estimates in that literature. However, as we emphasized earlier, the subsequent changes in the minimum wage take us well beyond the range of variation available for reduced-form minimum wage effects estimates. As we move across the columns, we see our estimates of the impacts of those further changes.⁵⁰ The ultimate impact once the \$15 minimum is fully phased in is given in the lower-right corner of the table and amounts to a reduction in the employment rate of 2.1 percentage points.

In Table 9, we investigate industrial variation in the impact of the minimum wage change on the wage structure in Seattle. In particular, we again raise each

⁴⁹ We believe that our estimates are best interpreted as providing an estimate of the effect of the minimum wage on the employment rate in a city, not on the employment level. A change in the minimum wage may well lead to migration of workers across cities, which is an additional difficulty in trying to evaluate effects of the minimum wage on employment levels. However, as we find that the effect of migration is to increase employment proportionally, we do not need to take a stance on how minimum wages affect intercity migration patterns when evaluating the effects on employment rates.

⁵⁰ Note that we use the wage and employment rate values predicted at the end of each year as the basis for the calculations in the following year.

TABLE 8—EMPLOYMENT IMPACTS FROM MINIMUM WAGE CHANGES: SEATTLE

	2015 (1)	2016 (2)	2017 (3)	2018 (4)	2019 (5)
1. Initial	71.2 (0.98)	70.7 (1.00)	70.2 (1.08)	69.6 (1.22)	69.3 (1.28)
2. Wage only	−1.40 (0.22)	−2.64 (0.43)	−3.98 (0.68)	−4.05 (0.85)	−4.55 (0.95)
3. Congestion + (2)	−0.44 (0.21)	−0.54 (0.25)	−0.62 (0.27)	−0.24 (0.10)	−0.25 (0.11)
4. End-of-year	70.7 (1.00)	70.2 (1.08)	69.6 (1.22)	69.3 (1.28)	69.1 (1.34)
Total employment change					−2.10
Standard error					0.93

Notes: Employment impacts calculated as follows: row 1 gives the employment rate at the beginning of the period. Row 2 gives the first round effect on the employment rate ($\beta_1 \times \Delta w_{ict}$, summed over industries). Row 3 gives the employment effect taking into account congestion effects $\left(\frac{\beta_1}{1 - \beta_3} \times \Delta w_{ct} \right)$. Row 4 gives the end-of-year employment rate.

sub-minimum wage worker to the minimum wage relevant for the given year and the worker's firm size but do so separately for each one-digit industry. These calculated wage effects vary with the initial year wage distribution in each industry and so, not surprisingly, the biggest wage effects are in the lowest-paid sector: the personal services and restaurant sector. There, the initial increases in the minimum wage imply over 6 percent increases in the average wage in each of the first three years. In contrast, in the financial sector, where the initial wage is 0.6 log points higher than in personal services, the effects on the average wage are always below 2 percent and often below 1 percent. In Table 10, we combine these wage changes with coefficients from a specification in which we allow β_1 and β_3 to vary by industry to obtain industry-specific employment rate effects. The effects presented here incorporate the congestion externality effects. The key point from this table is that our overall average employment rate effect from Table 8 conceals considerable variation by industry, with the employment rate declines in the service and restaurant industry being approximately triple those in the higher paying manufacturing sector.

In online Appendix B, we present the same analyses of the scheduled moves to \$15 minimum wages in Los Angeles and San Francisco. Table 11 summarizes those results along with the results for Seattle. For each city, we show the initial wage and employment rate and the final wage, employment rate, and percentage of workers whose wages are raised. The third column for each city shows the long-term changes in each outcome. The Los Angeles wage distribution is located considerably to the left of that for Seattle, resulting in a much larger proportion of workers directly affected by the minimum wage increase. Following from this, the average wage in Los Angeles is predicted to rise by 17 percent and the employment rate to fall by 3 percentage points: both substantially larger than for Seattle. San Francisco lies on the other side of Seattle, with a predicted average wage increase of 6 percent and a decline in its employment rate of 1.16 percentage points. Thus, San Francisco, with its higher wage distribution will face a smaller adjustment to a \$15 minimum wage than the other two cities.

TABLE 9—WAGE IMPACTS BY SECTOR FROM MINIMUM WAGE CHANGES: SEATTLE

	2014 (1)	2015 (2)	2016 (3)	2017 (4)	2018 (5)	2019 (6)	Total (7)
Agriculture, mining, cons.	3.10 (0.071)	0.020 (0.0062)	0.016 (0.0052)	0.018 (0.0047)	0.015 (0.0035)	0.016 (0.0035)	0.085 (0.020)
Manufacturing	3.20 (0.043)	0.0087 (0.0024)	0.013 (0.0028)	0.018 (0.0030)	0.0077 (0.0016)	0.0088 (0.0016)	0.057 (0.0092)
Transport, com., util.	3.10 (0.054)	0.016 (0.0043)	0.020 (0.0046)	0.024 (0.0044)	0.0097 (0.0024)	0.0099 (0.0023)	0.080 (0.015)
Retail, wholesale	2.87 (0.039)	0.025 (0.0034)	0.035 (0.0043)	0.041 (0.0042)	0.010 (0.0017)	0.010 (0.0017)	0.12 (0.013)
F.I.R.E.	3.19 (0.057)	0.0094 (0.0039)	0.014 (0.0041)	0.016 (0.0043)	0.0056 (0.0023)	0.0059 (0.0023)	0.051 (0.014)
Personal, entertainment	2.56 (0.032)	0.062 (0.0049)	0.064 (0.0047)	0.064 (0.0042)	0.027 (0.0027)	0.027 (0.0026)	0.24 (0.015)
Professional	3.15 (0.026)	0.011 (0.0015)	0.017 (0.0018)	0.021 (0.0018)	0.0082 (0.0010)	0.0093 (0.0010)	0.066 (0.0059)

Notes: Wage impacts by sector. Bootstrap standard errors in parentheses. The first column shows the wage in the sector at the onset of the policy. The rest of the columns show the change in the average log wage in the sector induced by the minimum wage policy implementation from one year to the next.

TABLE 10—EMPLOYMENT RATE BY SECTOR AND MINIMUM WAGE ROLL-OUT: SEATTLE

	2014 (1)	2015 (2)	2016 (3)	2017 (4)	2018 (5)	2019 (6)	Total (7)
Agriculture, mining, cons.	3.71 (0.42)	3.67 (0.42)	3.64 (0.42)	3.61 (0.42)	3.58 (0.42)	3.55 (0.42)	−0.16 (0.13)
Manufacturing	9.34 (0.59)	9.22 (0.58)	9.04 (0.58)	8.79 (0.59)	8.69 (0.59)	8.58 (0.60)	−0.77 (0.27)
Transport, com., util.	5.72 (0.53)	5.61 (0.53)	5.48 (0.53)	5.32 (0.54)	5.26 (0.55)	5.20 (0.55)	−0.52 (0.23)
Retail, wholesale	9.94 (0.64)	9.75 (0.63)	9.48 (0.63)	9.19 (0.64)	9.11 (0.64)	9.04 (0.65)	−0.91 (0.28)
F.I.R.E.	4.62 (0.44)	4.58 (0.44)	4.53 (0.44)	4.46 (0.44)	4.44 (0.43)	4.41 (0.43)	−0.21 (0.086)
Personal, entertainment	9.82 (0.66)	9.23 (0.64)	8.66 (0.66)	8.14 (0.69)	7.93 (0.71)	7.73 (0.73)	−2.09 (0.52)
Professional	28.0 (0.98)	27.7 (0.97)	27.3 (1.00)	26.7 (1.06)	26.5 (1.09)	26.3 (1.13)	−1.74 (0.65)

Notes: Employment rate by sector. Employment rate is calculated as as employment divided by working-age population. Each entry is the employment rate in an industry aggregate after policy year t . Bootstrap standard errors in parentheses.

Up to this point we have focused on the effects of minimum wage changes for the entire labor force. However, increases in the minimum wage are also likely to have distributional effects since initially low-paid workers will experience greater increases in wages than initially higher-paid workers. As noted in Section VII, if we treat individuals in different wage bins as workers with different skills, we can use our estimates to form predictions of how increases in the minimum wage affect different segments of the population. In particular, we predict the impact on employment rates of an increase in the minimum wage to \$15 for different groups defined

TABLE 11—SUMMARY OF MINIMUM WAGE RESULTS

	Seattle			San Francisco			Los Angeles		
	2015 (1)	2019 (2)	2019 – 2015 (3)	2015 (4)	2019 (5)	2019 – 2015 (6)	2015 (7)	2021 (8)	2021 – 2015 (9)
Employment	71.2 (0.98)	69.1 (1.34)	–2.10 (0.93)	69.1 (0.88)	67.8 (1.05)	–1.33 (0.63)	64.0 (0.49)	60.7 (1.24)	–3.32 (1.14)
Wage	3.04 (0.017)	3.13 (0.017)	0.098 (0.0045)	3.18 (0.017)	3.25 (0.017)	0.067 (0.0030)	2.85 (0.0096)	3.02 (0.0096)	0.17 (0.0032)
Fraction			0.30 (0.012)			0.25 (0.010)			0.42 (0.0070)

Notes: The first row contains the results for the employment rate, the second row contains the results for the average city log wage, and the third row contains the results for the fraction of workers impacted by the policy over the entire roll-out. The columns show the results for the start and end periods of the policy, along with the total change. Standard errors in parentheses, are constructed from a bootstrap of 500 replications.

as having initial wages below a given wage cutoff. For example, we can look at the effect of the minimum wage on all workers who are initially paid below \$15 hour or alternatively below \$10. We plot these distribution effects in Figure 1, reporting predicted changes in the employment rate for a set of thresholds defined by cutoffs at each dollar between \$10 and \$24 per hour (i.e., at \$10, \$11, \$12, . . . , \$24). There are three lines in the figure, one for each city we study. The bin marked 25 actually corresponds to the aggregate effect previously reported. As can be seen in this figure, for workers below \$10 per hour in Seattle, the employment rate declines by over 10 percent in response to raising the minimum wage to \$15. Meanwhile, for the larger group with wages at or below \$15, the decline is approximately 7 percent. In line with the aggregate effects, the negative employment implications of the minimum wage change are greatest for Seattle and smallest for San Francisco due to the relative initial locations of their wage distributions. Note that the effects report on this figure for low-wage workers are quite large as compared to the implied aggregate effects, indicating that such policies imply substantially different employment effects across the population.

VIII. Conclusion

In this paper, we present an empirically tractable labor demand framework which incorporates several insights from the macro-labor literature. The data we use to evaluate the framework involve city-industry level observations that span a period of four decades. Although our proposed labor demand framework is extremely parsimonious, we find considerable empirical support for it in the sense that (i) estimates of the main forces implied by the model are of the theoretically predicted sign and are statistically significant, (ii) over-identifying restrictions implied by the theory are not rejected by our data, and (iii) the results are robust and consistent across different levels of aggregation.

Our main motivation for reexploring the issue of labor demand was to shed light on the question: how does a reduction in the labor costs borne by firms affect the employment prospects of individuals. As noted in the introduction, there remains




FIGURE 1. PREDICTED DISTRIBUTIONAL EFFECTS OF MINIMUM WAGE CHANGES

Notes: Distribution impacts of the minimum wage changes are calculated by taking the average of $\frac{\beta_{li}}{1 - \beta_3} \times \max(0, 15 - w_q)$ for those workers with wage less than or equal to w_q . All figures are in 2014 dollars.

considerable debate over this question. Some researchers infer that labor demand is very elastic based on how economies react to migration flows while others infer that it is quite inelastic based on, for example, the observed effects of minimum wage changes. Our framework offers a reconciliation of these two views by separating out wage effects and population growth effects. Looking at the data through the lens of our model, we found there to be a significant negative effect of wages on employment, with an elasticity of close to -1 at the industry level and an elasticity of -0.28 at the city level. We argue that the lower elasticity at the city level is consistent with congestion externalities driven by search frictions. We also find that, holding wages constant, an increase in population is associated with a proportional increase in employment. We argue this latter pattern is consistent with the view that potential job creators are a special scarce factor because it is a scarce factor that is likely proportional to the population. An important insight we draw from our analysis is the importance of allowing a role for scarce entrepreneurial talent in the determination of labor demand.

REFERENCES

- Ahn, Tom, Peter Arcidiacono, and Walter Wessels. 2011. "The Distribution of Minimum Wage Increases When Labor Supply and Labor Demand Are Endogenous." *Journal of Business & Economic Studies* 29 (1): 12–23.
- Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6): 2121–68.
- Baker, Michael, Dwayne Benjamin, and Shuchita Stanger. 1999. "The Highs and Lows of the Minimum Wage Effect: A Time-Series Cross-Section Study of the Canadian Law." *Journal of Labor Economics* 17 (2): 318–50.

- Bartik, Timothy J.** 1991. *Who Benefits from State and Local Economic Policies?* Kalamazoo, MI: Upjohn Press.
- Bartik, Timothy J.** 1993. "Who Benefits from Local Job Growth: Migrants or the Original Residents?" *Regional Studies* 27 (4): 297–311.
- Bartik, Timothy J.** 2009. "How Do the Effects of Local Growth on Employment Rates Vary with Initial Labor Market Conditions?" W.E. Upjohn Institute for Employment Research Working Paper 09–148.
- Beaudry, Paul, David A. Green, and Benjamin Sand.** 2012. "Does Industrial Composition Matter for Wages? A Test of Search and Bargaining Theory." *Econometrica* 80 (3): 1063–104.
- Beaudry, Paul, David A. Green, and Benjamin M. Sand.** 2014. "Spatial Equilibrium with Unemployment and Wage Bargaining: Theory and Estimation." *Journal of Urban Economics* 79: 2–19.
- Beaudry, Paul, David A. Green, and Ben M. Sand.** 2018. "In Search of Labor Demand: Dataset." *American Economic Review*. <https://doi.org/10.1257/aer.20141374>.
- Belman, Dane, and Paul J. Wolfson.** 2014. "The New Minimum Wage Research." *Employment Research Newsletter* 21 (2): 4–5.
- Blanchard, Olivier Jean, and Lawrence F. Katz.** 1992. "Regional Evolutions." *Brookings Papers on Economic Activity* (1): 1–61.
- Blau, Francine D., and Lawrence M. Kahn.** 1999. "Institutions and Laws in the Labor Market." In *Handbook of Labor Economics*, Vol. 3, edited by Orley Ashenfelter and David Card, 1399–1461. Amsterdam: Elsevier.
- Bound, John, and Harry J. Holzer.** 2000. "Demand Shifts, Population Adjustments, and Labor Market Outcomes during the 1980s." *Journal of Labor Economics* 18 (1): 20–54.
- Brochu, Pierre, and David A. Green.** 2013. "The Impact of Minimum Wages on Labour Market Transitions." *Economic Journal* 123 (573): 1203–35.
- Brown, James N., and Orley Ashenfelter.** 1986. "Testing the Efficiency of Employment Contracts." *Journal of Political Economy* 94 (3): S40–S87.
- Cahuc, Pierre, Stéphane Carcillo, and Thomas Le Barbanchon.** 2014. "Do Hiring Credits Work in Recessions? Evidence from France." Institute for the Study of Labor (IZA) Discussion Papers 8830.
- Card, David.** 1990a. "The Impact of the Mariel Boatlift on the Miami Labor Market." *Industrial and Labor Relations Review* 43 (2): 245–57.
- Card, David.** 1990b. "Unexpected Inflation, Real Wages, and Employment Determination in Union Contracts." *American Economic Review* 80 (4): 669–88.
- Card, David.** 2001. "Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration." *Journal of Labor Economics* 19 (1): 22–64.
- Cooper, Russell, John Haltiwanger, and Jonathon L. Willis.** 2015. "Dynamics of Labor Demand: Evidence from Plant-Level Observations and Aggregate Implications." *Research in Economics* 69 (1): 37–50.
- Dahl, Gordon B.** 2002. "Mobility and the Return to Education: Testing a Roy Model with Multiple Markets." *Econometrica* 70 (6): 2367–420.
- Dickens, Richard, and Alan Manning.** 2004. "Has the National Minimum Wage Reduced UK Wage Inequality?" *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 167 (4): 613–26.
- Evans, David, and Ben Zipperer.** 2014. *Where Does Your State's Minimum Wage Rank against the Median Wage?* Washington, DC: Washington Center for Equitable Growth.
- Flinn, Christopher J.** 2006. "Minimum Wage Effects on Labor Market Outcomes under Search, Matching, and Endogenous Contact Rates." *Econometrica* 74 (4): 1013–62.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift.** 2017. "Bartik Instruments: What, When, How, and Why." NBER Working Paper 24408.
- Hamermesh, Daniel.** 1993. *Labor Demand*. Princeton, NJ: Princeton University Press.
- Hopenhayn, Hugo A.** 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (5): 1127–50.
- Jensen, J. Bradford, and Lori G. Kletzer.** 2006. "Trade Services: Understanding the Scope and Impact of Services Outsourcing." In *Brookings Trade Forum, Offshoring White-Collar Work*, Vol. 6, edited by Susan M. Collins and Lael Brainard, 75–134. Washington, DC: Brookings Institution.
- King, Robert G., and Julia K. Thomas.** 2006. "Partial Adjustment without Apology." *International Economic Review* 47 (3): 779–809.
- Kramarz, Francis, and Marie-Laure Michaud.** 2010. "The Shape of Hiring and Separation Costs in France." *Labour Economics* 17 (1): 27–37.
- Kramarz, Francis, and Thomas Philippon.** 2001. "The Impact of Differential Payroll Tax Subsidies on Minimum Wage Employment." *Journal of Public Economics* 82 (1): 115–46.

- Krueger, Alan B., and Jörn-Steffen Pischke.** 1997. "Observations and Conjectures on the U.S. Employment Miracle." Princeton University, Department of Economics, Industrial Relations Section Working Papers 769.
- Kuhn, Peter.** 1988. "Unions in a General Equilibrium Model of Firm Formation." *Journal of Labor Economics* 6 (1): 62–82.
- Levine, Ross, and Yona Rubinstein.** 2013. "Smart and Illicit: Who Becomes an Entrepreneur and Does It Pay?" NBER Working Papers 19276.
- Light, Audrey, and Robert Munk.** 2015. "Business Ownership vs. Self-Employment." Unpublished.
- MaCurdy, Thomas E., and John H. Pencavel.** 1986. "Testing between Competing Models of Wage and Employment Determination in Unionized Markets." *Journal of Political Economy* 94 (S3): S3–S39.
- Manning, Alan, and Michael Amior.** 2015. "The Persistence of Local Joblessness." London School of Economics and Political Science Centre for Economic Performance Discussion Papers 1357.
- Meer, Jonathan, and Jeremy West.** 2016. "Effects of the Minimum Wage on Employment Dynamics." *Journal of Human Resources* 51 (2): 500–522.
- Mutti, John, and Shelby Gerking.** 1983. "Changes in Income Distribution and Welfare from Greater Immigration of Unskilled Workers." *European Economic Review* 23 (1): 103–16.
- Peri, Giovanni, and Chad Sparber.** 2009. "Task Specialization, Immigration, and Wages." *American Economic Journal: Applied Economics* 1 (3): 135–69.
- Reich, Michael, Ken Jacobs, Annette Bernhardt, and Ian Perry.** 2014. *The Impact of Oakland's Proposed City Minimum Wage Law: A Prospective Study*. Berkley, CA: Center on Wage and Employment Policy Brief, University of California Berkley.
- Ruggles, Steven, Trent J. Alexander, Katie Genadek, Ronald Goeken, Matthew Schroader, and Matthew B. Sebek.** 2015. Integrated Public Use of Microdata Series: Version 6.0 [Machine-readable database]. Minneapolis, MN: Minnesota's Population Center.
- Sorkin, Isaac.** 2015. "Are There Long-Run Effects of the Minimum Wage?" *Review of Economic Dynamics* 18 (2): 306–33.
- Tschopp, Jeanne.** 2015. "The Wage Response to Shocks: The Role of Inter-Occupational Labour Adjustment." *Labour Economics* 37: 28–37.
- Tschopp, Jeanne.** 2017. "Wage Formation: Towards Isolating Search and Bargaining Effects from the Marginal Product." *Economic Journal* 127 (603): 1693–729.
- Wooldridge, Jeffrey M.** 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.

This article has been cited by:

1. Jianhua Sun, Shaobo Hou, Yuxia Deng, Huaicheng Li. 2024. New media environment, green technological innovation and corporate productivity: Evidence from listed companies in China. *Energy Economics* **131**, 107395. [[Crossref](#)]
2. Tristan Potter. 2024. Destabilizing search technology. *Journal of Monetary Economics* **90**, 103557. [[Crossref](#)]
3. Luis Medrano-Adán, Vicente Salas-Fumás. 2023. Do minimum wages deliver what they promise? Effects of minimum wage on employment, output, and income inequality from occupational choice theory. *Economic Analysis and Policy* **80**, 366-383. [[Crossref](#)]
4. João Galindo da Fonseca, Charles Berubé. 2023. SPOUSES, CHILDREN, AND ENTREPRENEURSHIP. *International Economic Review* **64**:3, 1165-1199. [[Crossref](#)]
5. Bo Chen, Dong Tan. 2023. Industrial Robots and the Employment Quality of Migrant Workers in the Manufacturing Industry. *Sustainability* **15**:10, 7998. [[Crossref](#)]
6. Wenjian Xu. 2023. Employment Decline During the Great Recession: the Role of Firm Size Distribution. *The Economic Journal* **133**:652, 1586-1625. [[Crossref](#)]
7. Massimo Anelli, Gaetano Basso, Giuseppe Ippedico, Giovanni Peri. 2023. Emigration and Entrepreneurial Drain. *American Economic Journal: Applied Economics* **15**:2, 218-252. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
8. Gert Bijlens, Shyngys Karimov, Jozef Konings. 2023. Does Automatic Wage Indexation Destroy Jobs? A Machine Learning Approach. *De Economist* **171**:1, 85-117. [[Crossref](#)]
9. DOMENICO FERRARO. 2023. Fast Rises, Slow Declines: Asymmetric Unemployment Dynamics with Matching Frictions. *Journal of Money, Credit and Banking* **55**:2-3, 349-378. [[Crossref](#)]
10. Tim Bartik. 2023. Seize the Time: Needed Research on Local Economic Development in an Era of Increased Attention to Problems of Place. *Economic Development Quarterly* **37**:1, 7-13. [[Crossref](#)]
11. David A. Green, Huju Liu, Yuri Ostrovsky, Garnett Picot. 2023. Are Immigrants Particularly Entrepreneurial? Policy Lessons from a Selective Immigration System. *SSRN Electronic Journal* **4**. . [[Crossref](#)]
12. Vikrant Kamble, Dylan Brewer. 2023. The effect of mass layoffs on related industries: Evidence from a mining ban in India. *SSRN Electronic Journal* **138**. . [[Crossref](#)]
13. Wenjian Xu. 2022. Social insurance and entrepreneurship: The effect of unemployment benefits on new-business formation. *Strategic Entrepreneurship Journal* **16**:3, 522-551. [[Crossref](#)]
14. Wenbo Zhu. 2022. Hollowing out and slowing growth: The role of process innovations. *Review of Economic Dynamics* **45**, 217-236. [[Crossref](#)]
15. João Galindo da Fonseca. 2022. Unemployment, entrepreneurship and firm outcomes. *Review of Economic Dynamics* **45**, 322-338. [[Crossref](#)]
16. Baxtiyor Ganiev. 2022. The state of demand for labor in the private sector in Uzbekistan and its importance in providing employment. *Economics and Innovative Technologies* **10**:3, 45-54. [[Crossref](#)]
17. Gaetano Perone. 2022. The effect of labor market institutions and macroeconomic variables on aggregate unemployment in 1990–2019: Evidence from 22 European countries. *Industrial and Corporate Change* **31**:2, 500-551. [[Crossref](#)]
18. Michael Lokshin, Martin Ravallion. Would Mexican Migrants be Willing to Guarantee Americans a Basic Income? . [[Crossref](#)]
19. Maria Iosifidi, Ekaterini Panopoulou, Chris Tsoumas. 2021. Mortgage loan demand and banks' operational efficiency. *Journal of Financial Stability* **53**, 100851. [[Crossref](#)]

20. Arnab K. Basu, Nancy H. Chau, Vidhya Soundararajan. 2021. Contract employment as a worker discipline device. *Journal of Development Economics* **149**, 102601. [[Crossref](#)]
21. Andrew C. Johnston. 2021. Unemployment Insurance Taxes and Labor Demand: Quasi-Experimental Evidence from Administrative Data. *American Economic Journal: Economic Policy* **13**:1, 266-293. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
22. Michael Joffe. Equilibrium, Instability, Growth and Feedback in Economics 43-68. [[Crossref](#)]
23. Pierre Deschamps, Guillaume Wilemme. 2021. Regional Unemployment Persistence and Agglomeration Effects. *SSRN Electronic Journal* **108**. . [[Crossref](#)]
24. Greg Howard. 2020. The Migration Accelerator: Labor Mobility, Housing, and Demand. *American Economic Journal: Macroeconomics* **12**:4, 147-179. [[Abstract](#)] [[View PDF article](#)] [[PDF with links](#)]
25. Jorge Pérez Pérez. 2020. The minimum wage in formal and informal sectors: Evidence from an inflation shock. *World Development* **133**, 104999. [[Crossref](#)]
26. Mario Bossler, Michael Oberfichtner, Claus Schnabel. 2020. Employment Adjustments Following Rises and Reductions in Minimum Wages: New Insights From a Survey Experiment. *LABOUR* **34**:3, 323-346. [[Crossref](#)]
27. German P. Pupato, Ben Sand, Jeanne Tschoopp. 2019. Estimating the Gains From Trade in Frictional Labor Markets. *SSRN Electronic Journal* **34**. . [[Crossref](#)]
28. Wenjian Xu. 2019. Entrepreneurship Among the Unemployed: The Effects of Unemployment Benefits. *SSRN Electronic Journal* **117**. . [[Crossref](#)]