Libraries and Learning Services

University of Auckland Research Repository, ResearchSpace

Suggested Reference

Greenaway-McGrevy, R., & Hood, K. (2016). *Worker migration or job creation? Persistent shocks and regional recoveries* (285). Auckland: University of Auckland.

Copyright

Items in ResearchSpace are protected by copyright, with all rights reserved, unless otherwise indicated. Previously published items are made available in accordance with the copyright policy of the publisher.

<https://researchspace.auckland.ac.nz/docs/uoa-docs/rights.htm>

WORKER MIGRATION OR JOB CREATION? PERSISTENT SHOCKS AND REGIONAL RECOVERIES

By

Ryan Greenaway-McGrevy and Kyle Hood

March 2016

DEPARTMENT OF ECONOMICS WORKING PAPER NO. 285

**THE UNIVERSITY OF AUCKLAND
NEW ZEALAND**

**DEPARTMENT OF ECONOMICS
UNIVERSITY OF AUCKLAND**

Private Bag 92019

Auckland, New Zealand

<http://econ.auckland.ac.nz>

Worker migration or job creation? Persistent shocks and regional recoveries

Ryan Greenaway-McGrevy* and Kyle Hood^{†‡}

March 14, 2016

Abstract

Although a large body of literature has documented the role of household out-migration in the recovery from regional downturns, the role that firms play in the recovery process has remained a neglected topic of research. Firms may choose to locate new jobs in depressed regions, thereby reducing unemployment through the job creation channel. We present a new empirical model of regional adjustment that permits us to decompose recoveries into both household and firm responses to local economic conditions. The model features a set of auxiliary serial dependence parameters that are used to filter out persistency in the identified labor market shocks, so that changes in employment obtained from the fitted model only reflect the endogenous firm response of interest, and not the ongoing exogenous job destruction from the original downturn. We find that the labor demand response is two to three times larger than the labor supply response, meaning that local job creation—and not household out-migration—is the main driver of recoveries in the US. This result is robust to a wide variety of model specifications and identification strategies.

Keywords: Regional equilibration, Labor mobility, Job creation, Migration.

JEL classification codes: J61, R23, R30, C33

*Department of Economics, The University of Auckland, Auckland, New Zealand. Corresponding Author. Email: r.mcgrevy@auckland.ac.nz. Address: The University of Auckland Business School, 12 Grafton Road, Private Bag 92019, Auckland 1142, New Zealand.

[†]Bureau of Economic Analysis, Washington, DC.

[‡]We would like to thank Benjamin Bridgman, Steffen Lippert, Erika McEntarfer and participants at the 25th meeting of the NZESG, the 11th International WEAI Conference, the AREUEA National Conference 2015, and the University of North Carolina Charlotte Department of Economics Seminar Series for their invaluable comments. Any remaining errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the Bureau of Economic Analysis or the Department of Commerce.

1 Introduction

Geographic regions of an economy are periodically subjected to adverse shocks. Although a large literature has examined how such shocks are absorbed through out-migration, unemployment and labor force participation,¹ the role that firms play in mitigating the effect of these shocks over the long run has remained a neglected topic of research. For example, while households may relocate to more prosperous regions of the economy, firms may also create new jobs in depressed regions, attracted by the surplus of labor and low wages. The job creation generated by this labor demand response therefore represents a parallel channel to out-migration through which the labor market can recover.

In this paper we provide a novel approach for measuring the contribution of the endogenous labor demand response to regional recoveries. Our approach builds on the structural vector autoregressions (VARs) that are often used to ascribe variation in regional employment, labor force and population levels to structural labor market shocks (see Blanchard and Katz, 1992; Decressin and Fatás, 1995; Tani, 2003; amongst others). We begin by showing that persistency in the identified structural shocks is the primary impediment to measuring the effects of the endogenous labor demand response in these empirical models. Through the lens of the underlying economic model, changes in regional employment levels embody labor market shocks as well as the subsequent household and firm responses to those shocks. If the structural shocks identified in the empirical model are persistent—in the sense that the economic shock occurs over several time periods—it is difficult to disentangle the increase in employment due to the endogenous labor demand response from the ongoing decrease in employment due to the original exogenous downturn.

In order to estimate the effect of the endogenous labor demand response we suggest modelling the response to a one-period labor demand shock (i.e., a shock that exhibits no serial dependence). We show that this can be achieved by parameterizing serial dependence into the empirical model for the purposes of model estimation (as is also implicitly done in the SVAR literature), but “switching off” these auxiliary serial dependence parameters when using the fitted model to impute responses

¹See, amongst others, Marston (1985), Topel (1986), Blanchard and Katz (1992) and Bartik (1991,1993).

to labor market shocks. This means that subsequent changes in employment and population exclusively reflect the supply and demand responses of labor and firms—because these changes are not contaminated with persistence in the shock. Using this approach we find that the labor demand response is two to three times larger than the labor supply response, meaning that local job creation—and not household out-migration—drives the recovery. The recovery is however highly protracted, extending over a period in excess of twenty years. Taken together, these findings suggest that economic shocks can have large, long-lasting effects on local residents.

Our main finding of a highly-protracted, largely firm-mediated adjustment process contrasts with early work by Marston (1985) and Blanchard and Katz (1992), who conclude that there is a comparatively large and rapid migration response to regional shocks. It is however consistent with a broader literature which concludes that migration plays a relatively small role in equilibrating regional unemployment rates, and that labor market shocks can have long-lasting effects on regional labor markets (see the discussion in Partridge *et al.*, 2015). Our findings are of particular interest because our methodology is in the macroeconometric tradition of Blanchard and Katz (1992), but due to our careful treatment of shock persistency in the empirical model, our conclusions bear more similarity to the consensus.

The economic model of regional adjustment employed by Blanchard and Katz (1992) (which will be subsequently referred to as the “BK model”) provides the foundation of our analysis, and has been used in other contexts to model migration in response to regional labor market disparities (Decressin and Fatás, 1995; Debelle and Vickery, 1999; Tani, 2003; Grimes, Maré and Morton, 2009). The model is often empirically specified as a VAR (see Blanchard and Katz, 1992; Decressin and Fatás, 1995; Tani, 2003; Saks, 2008; Zabel, 2012, amongst others). However, a shortcoming of the Blanchard and Katz (1992) VAR in its present application is that the labor demand response can only be estimated under a rather restrictive condition on the structural shocks—they must be uncorrelated over time. Because the shocks identified in these models tend to in fact exhibit significant (positive) serial dependence, inferring the labor demand response based on changes in employment levels will lead us to understate the role of firms in the recovery—and to correspondingly overstate the contribution of migration to the recovery. We illustrate weakness below in more detail

in section 2 below.

To overcome this problem we start by specifying the BK model as a structural vector error correction model (VECM). Within the VECM the structural parameters governing the firm and worker responses are separate from the auxiliary parameters that capture any persistence in the underlying shocks. This means that the VECM presents a simple framework for untangling the endogenous responses of interest from the ongoing exogenous variation due to persistency in the shock. We can then obtain clean estimates of the endogenous responses by “switching off” the serial dependence parameters when imputing impulse responses to the shocks of interest.

One of the main contributions of the paper is therefore to provide a simple generalization of the BK model that explicitly addresses with the problems caused by persistent shocks in the empirical model. Indeed its inability to accommodate various forms of persistence has remained a widely perceived weakness of the original Blanchard and Katz (1992) method. Hall (1992) originally pointed out that we must assume that labor demand shocks are uncorrelated over time if we want to obtain the endogenous firm response from the fitted model,² while Bartik (1993) and Obstfeld and Peri (1998) question the assumption in the BK model that labor demand shocks cannot have a permanent effect on regional unemployment and participation rates. Three main pieces of evidence establish that the identified labor demand shocks exhibit persistence. The first is based on the original Blanchard and Katz (1992) and related results, in which labor demand shocks clearly exhibit inertia. For example, after an initial negative shock, employment continues to fall for several periods before eventually beginning to recover. Because there is no mechanism within the BK model to generate this observed behavior, the inertia can only be rationalized as positive serial dependence in the initial shock (without persistency, employment must *immediately* start moving in the *opposite* direction of the shock, to “correct” its initial effect). Second, a large literature shows that regional employment rates are highly persistent (even after conditioning on a national average), and are often indistinguishable from unit root processes (see the survey in Partridge *et al.*, 2015). In order to accommodate this persistency within the BK model we can permit positive

²See p. 65 of Blanchard and Katz (1992), where Hall states that “They seem to assume that demand shifts are inherently one-time random-walk shifts, so that the effects a year later and after are the result of endogenous responses. But no basis exists for this implicit assumption. The demand shift itself could have any form of persistence.”

serial dependence in the underlying structural shocks. The third piece of evidence we introduce is that at the industry level, local industry-specific shocks have delayed effects on other industries (for example, construction, and trade, transportation, and utilities). This strongly suggests that industrial input-output relationships are generating the observed persistence in labor demand shocks identified based on regional aggregates.

Our findings hold up under an array of model specifications, estimation methods, identification strategies and sample periods. Our baseline empirical model is based on a data set constructed from metropolitan statistical area (MSA) employment, labor force and working age population over the 1990 to 2012 period. We use MSA data because an MSA is (by definition) an approximation to a single regional labor market. Other specifications include a longer time series of MSA data (1970 to 2012), a data set that includes a direct measure of net migration, and model specifications that account for household migration frictions such as wages, housing prices and housing supply elasticities. We also estimate the baseline model using an instrumental variables estimator in order to circumvent potential bias in the OLS estimator. To identify labor demand shocks we use the Bartik (1991) shift-share instrument; however we also employ recursive identification (as used by Blanchard and Katz, 1992) as an additional robustness check. All cases present a nearly identical picture, with the labor demand response playing a larger role than out-migration in a highly-protracted adjustment process. A simple extension of the empirical model indicates that the number of establishments in the MSA is relatively constant during the recovery, suggesting that the increase in regional labor demand is mainly due to existing firms expanding their payrolls.

We organize the remainder of the paper as follows: We conclude the introduction with a brief literature review. The following section outlines the problems caused by serial dependence in structural shocks. After this, we discuss how using an empirical VECM model allows us to isolate the effects of serial dependence in the empirical model. Section 3 discusses the data used in the model. Our empirical results are contained in section 4. We then consider extensions and robustness checks of our baseline model. In section 6 we analyze the underlying mechanism(s) generating shock persistence, and present evidence suggesting that local input-output or multiplier effects are a substantial driver of shock persistency. Finally, we conclude the paper with a detailed discussion

of the implications of our finding

Related literature This work relates to a number of literatures that focus on how regional economies absorb and recover from labor demand shocks, including (but not limited to) Marston (1985), Topel (1986), Greenwood *et al.* (1986), Blanchard and Katz (1992), Bartik (1991, 1993), and Lkhagvasuren (2012). The Blanchard and Katz (1992) approach has been replicated by Decressin and Fatás (1995) for the EU, who compare US and EU labor mobility. Debelle and Vickery (1999) apply the model to Australia while Grimes, Maré and Morton (2009) apply the model to New Zealand. Our work builds directly off of this literature.

A line of research has taken up the related question of whether amenities or job opportunities are the main driver of household migration (Greenwood and Hunt, 1989; Partridge and Rickman, 1999, 2003; and Chen and Rosenthal, 2008). This research focuses on decomposing household migration into labor supply and demand forces, whereas our focus is decomposing recoveries from shifts only in labor demand into household labor supply and firm labor demand responses. These two lines of research complement each other, as a holistic picture of migration of jobs and workers across regions would include the effects of both labor supply and demand shocks and the responses thereto.

Recent work has focussed on the role that housing plays in migration decisions. Glaeser and Gyourko (2005), Glaeser *et al.* (2006), Saiz (2007), Saks (2008), and Zabel (2012) explore the interplay between housing markets and migration, and much of this literature employs a VAR methodology. While our work does not focus on housing *per se*, our methods are applicable to specifications that are used in this literature, and we estimate models that include house prices and house price elasticities as extensions to our main empirical specification. In addition, the basic methodology that we suggest for modifying the panel VAR approach to studying house prices and migration could be combined with models that consider cross-sectional characteristics of regional house price dynamics, such as spatiotemporal dependence (Bailey, Holly and Pesaran, 2016).

2 Model

Our empirical approach builds on the canonical BK model of regional labor market dynamics. The BK model is often used to describe regional adjustments: It is used to infer labor mobility in other countries (Decressin and Fatás, 1995; Debelle and Vickery, 1999; Tani, 2003; Grimes, Maré and Morton, 2009) and to build models of geographic relocation in other contexts (see, amongst others, Borjas *et al.*, 1997; Borjas, 1999, 2006; Kline, 2008). In the model, workers (firms) are attracted to (repelled from) labor markets with relatively low unemployment and high wages, and repelled from (attracted to) labor markets with relatively high unemployment and low wages. These endogenous labor demand and supply responses thereby equilibrate the incidence of labor market shocks across different geographic regions. In this way, labor demand (and supply) shocks may have permanent effects on the local level of employment, but (real) wages, unemployment rates, and participation rates will be stationary after conditioning on a national average.

We present a brief summary of the BK model. All variables in the model are defined in logs and in relative terms. For example, $w_{i,t}$ is the log relative wage rate in location i at time t : If $W_{i,t}$ denotes the nominal wage and W_t denotes the national average wage, then $w_{i,t} = \log(W_{i,t}/W_t)$. Labor supply ($e_{i,t}^S$) and demand ($e_{i,t}^D$) in location i at time t are given by

$$e_{i,t}^S = \alpha_{i,t}^S + \gamma^S w_{i,t}, \quad e_{i,t}^D = \alpha_{i,t}^D - \gamma^D w_{i,t}. \quad (1)$$

Labor demand and supply may shift as follows

$$\Delta \alpha_{i,t}^S = \lambda^S w_{i,t-1} + \varepsilon_{i,t}^S, \quad \Delta \alpha_{i,t}^D = -\lambda^D w_{i,t-1} + \varepsilon_{i,t}^D, \quad (2)$$

where $\lambda^S > 0$ captures workers' endogenous labor supply response to local economic conditions (summarized by relative wages), $\lambda^D > 0$ captures firms' endogenous labor demand response, and $\varepsilon_{i,t}^S$ and $\varepsilon_{i,t}^D$ are the exogenous labor supply and demand shocks.³

³Note that we omit locational amenities from the endogenous response functions given in (2). These are often accounted for through location-specific parameters (see (5) and (6) in Blanchard and Katz, 1992). Including these parameters would have no material effect on the analysis in this section, but in our empirical specification we include cross-sectional fixed effects to account for these forms of locational amenities. Note that the BK model is focused

In the empirical analysis to follow we focus on measuring (a) the effect on employment of the endogenous labor demand response to a labor demand shock, and (b) the effect on working-age population of the endogenous labor supply response to a labor demand shock. Combining equations (1) and (2), solving for equilibrium employment (denoted $e_{i,t}$) we can obtain equations describing the evolution of wages and employment in reduced form as follows:

$$\begin{bmatrix} \Delta e_{i,t} \\ \Delta w_{i,t} \end{bmatrix} = \begin{bmatrix} \frac{\gamma^S}{\gamma^S + \gamma^D} & \frac{\gamma^D}{\gamma^S + \gamma^D} \\ \frac{1}{\gamma^S + \gamma^D} & -\frac{1}{\gamma^S + \gamma^D} \end{bmatrix} \begin{bmatrix} 0 & \lambda^D \\ 0 & -\lambda^S \end{bmatrix} \begin{bmatrix} e_{i,t-1} \\ w_{i,t-1} \end{bmatrix} + \begin{bmatrix} \frac{\gamma^S}{\gamma^S + \gamma^D} & \frac{\gamma^D}{\gamma^S + \gamma^D} \\ \frac{1}{\gamma^S + \gamma^D} & -\frac{1}{\gamma^S + \gamma^D} \end{bmatrix} \begin{bmatrix} \varepsilon_{i,t}^D \\ \varepsilon_{i,t}^S \end{bmatrix}. \quad (3)$$

Let $p_{i,t}$ denote the working age population. Following BK we replace $w_{i,t}$ with the broad employment rate $e_{i,t} - p_{i,t}$ by positing a negative linear relationship between wages and unemployment, namely $w_{i,t} = \theta(e_{i,t} - p_{i,t})$ for some $\theta > 0$. (Although BK make no distinction between involuntary unemployment and participation in the model, this condition is relaxed in the empirical model.)

With some algebra, we can then express the reduced-form BK model succinctly as follows:

$$\Delta Y_{i,t} = \mathbf{\Gamma} \mathbf{\Lambda} \boldsymbol{\delta}' Y_{i,t-1} + \mathbf{\Gamma} \varepsilon_{i,t}, \quad (4)$$

where

$$Y_{i,t} \equiv \begin{bmatrix} e_{i,t} \\ p_{i,t} \end{bmatrix}, \quad \mathbf{\Gamma} \equiv \begin{bmatrix} \frac{\gamma^S \theta}{\gamma^S + \gamma^D} & \frac{\gamma^D \theta}{\gamma^S + \gamma^D} \\ \frac{\gamma^S \theta - 1}{\gamma^S + \gamma^D} & \frac{\gamma^D \theta + 1}{\gamma^S + \gamma^D} \end{bmatrix}, \quad \mathbf{\Lambda} \equiv \begin{bmatrix} \lambda^D \\ -\lambda^S \end{bmatrix}, \quad \boldsymbol{\delta} \equiv \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \quad \varepsilon_{i,t} \equiv \frac{1}{\theta} \begin{bmatrix} \varepsilon_{i,t}^D \\ \varepsilon_{i,t}^S \end{bmatrix},$$

with the parameter vector $\mathbf{\Lambda}$ capturing the endogenous responses of interest. Blanchard and Katz (1992) make the explicit assumption that the labor market shocks $\varepsilon_{i,t}$ are white noise (see (5) and (6) in Blanchard and Katz, 1992). Under this assumption the economic model thus implies that growth in labor market variables $\Delta Y_{i,t}$ is described by a structural VECM of the form of (4). Note that under this assumption we could have just as easily expressed the BK model as a reduced-form VAR(1) in the vector $(\Delta e_{i,t}, p_{i,t} - e_{i,t})'$.

on the endogenous responses to labor market shocks, and not long-run convergence to spatial equilibrium based on regional amenities (see, e.g., Partridge et al, 2012; 2015), so that omission of these parameters for instructive purposes is warranted.

2.1 Shock Persistence

The BK model described (1) and (2) is often formulated empirically as a VAR when it is used to model regional equilibration in practice (Blanchard and Katz, 1992; Decressin and Fatás, 1995; Debelle and Vickery, 1999; Tani, 2003; and Grimes, Maré and Morton, 2009). Under suitable restrictions to identify regional labor demand shocks, the subsequent changes in working-age population, labor force and employment implied by the fitted model tell us about the incidence of and recovery from the shock. The responses of these labor market variables are typically interpreted through the lens of the economic model: Increases in employment after a negative shock are due to the endogenous labor demand response, while decreases in population are due to the endogenous labor supply response (out-migration of households).

Although the economic model in its strictest form implies a VAR(1), additional lags of the endogenous variables are often included to ensure that the fitted residuals are uncorrelated over time. For example, Blanchard and Katz (1992), Saks (2008) and Zabel (2012) all impose two lags in their empirical VARs. The most straightforward way to accommodate this feature of the data in the economic model is to relax the rather stringent assumption that the structural shocks from the economic model (4) are uncorrelated over time. However, once we allow the reduced-form errors in the economic model to exhibit persistence, we change the interpretation of the empirical model.⁴

We illustrate this point by taking a specific version of the reduced-form model given by (4) and feeding it shocks that are potentially serially dependent. We set $\lambda^D = -\lambda^S = 0.1$, so that firms and households contribute equally to the recovery, and we set $\gamma^D = 0$ so that short-run labor demand is perfectly inelastic. We first consider the case of a -1% labor demand shock that is serially uncorrelated. Figure 1 exhibits the paths of employment and population to the shock. There is an immediate reduction in employment, leading to an increase in the unemployment rate $p_{i,t} - e_{i,t}$ (the difference between the dotted and solid lines). From period two onwards the employment level begins to recover as (short run) labor demand increases in response to the high unemployment

⁴Serial dependence in the structural shocks of the economic model may be rationalized in terms of mechanisms that are omitted from the economic model. For example, Blanchard and Katz (1992, p. 23) discuss how capital accumulation and divestment could amplify labor market shocks over the short run. They do not, however, discuss the implications of this extension when using the empirical model to make statements about labor mobility.

Figure 1: Illustrative paths of log employment (e) and log population (p) after a -1% labor demand shock exhibiting no serial dependence. Worker and firm responses each account for half of the recovery.

rate (and, implicitly, low real wages). Meanwhile population decreases from period two onwards as (short run) labor supply decreases in response to the high unemployment rate. In the long run, there is a permanent -0.5% decrease in employment and population, reflecting the fact that firms and households contribute equally to the recovery.

Next we consider what a model with the same values for λ^S and λ^D (i.e., $\lambda^D = -\lambda^S = 0.1$) implies for employment and population in response to a serially dependent labor demand shock. Again, the initial shock will be -1%, but it will follow an AR(1) process with a coefficient of 0.5. Thus the initial shock will decay over time at a geometric rate, causing a further -0.5% loss of jobs in period 2, -0.25% in period 3, and so forth. Figure 2 exhibits the paths of employment and population in response to the shock. The initial shock causes employment to decrease in the first period by -1%. In the subsequent periods the change in the employment level reflects both the inertia in the initial shock as well as the effects of the endogenous labor demand response. In

Figure 2: Illustrative paths of log employment (e) and log population (p) after a -1% labor demand shock exhibiting moderate serial dependence ($AR(1) = 0.5$). Worker and firm responses each account for half of the recovery.

the first period after the shock (period 2), the former outweighs the latter, and the net change in employment levels is negative. As time progresses, the initial dependence in the shock continues to dissipate and so eventually the labor demand response is larger than the recurring echo of the shock. In this example, such a “tipping point” occurs three periods after the shock. Labor demand continues to increase, and labor supply continues to decrease, until the employment-to-population ratio returns to equilibrium.

The initial -1% labor demand shock leads to a 1% reduction in the long-run employment level. A naïve interpretation would be that worker movement must account for all of the recovery (or, equivalently, that there is no labor demand response to the shock) because there is no recovery in employment relative to the initial shock. However, we know this interpretation is incorrect, since we set the model parameters so that the endogenous labor demand response is exactly the same size as the labor supply response. In the absence of the labor demand response, the initial 1% fall in labor demand would ultimately result in a permanent $2\% = (1\% \times 1/(1 - 0.5))$ reduction in employment

Figure 3: Illustrative paths of log employment (e) and log population (p) after a -1% labor demand shock exhibiting very high serial dependence ($AR(1) = 0.9$). Worker and firm responses each account for half of the recovery.

(where recall that the $AR(1)$ parameter on the shock is 0.5). The fact that employment ends up falling by only 1% in the long run in our example—and not 2%—is because of the endogenous labor demand response. In addition, we can see that employment increases from period three onwards, indicating a net increase in labor demand. Most importantly, we cannot quantify the amount of employment generated by the endogenous labor demand response by simply examining the long-run effect of the shock on employment, and thus we actually cannot make statements about whether firms or workers contribute more to the recovery based on this figure.

If the labor demand shock is very persistent we may not observe a recovery in employment at all. Figure 3 exhibits the responses when the $AR(1)$ parameter is 0.9.

The impulse responses documented in the empirical literature often resemble Figure 2 more than Figures 1 or 3. The employment response exhibits an inverse “hump” shape, indicating that the initial labor demand shock is amplified in the short run before the employment level begins to increase (note this becomes a standard hump shape in the case of a positive shock). Figure 4 in

the empirical section of the paper as well as the main results of Blanchard and Katz (1992) (see fig. 7, p. 33) clearly exhibit the initial amplification followed by a recovery. Similar patterns can be observed in Figure 2 of Saks (2008), Figure 1 of Partridge and Rickman (2012) and Figure 1 of Partridge and Rickman (2006). (Note however that the recovery in the case of the latter two examples is rather limited.)

These “hump-shaped” employment responses are often interpreted as implying that the household migration is the primary driver of recoveries. For example, Blanchard and Katz (1992) state that “a state typically returns to normal after an adverse shock not because employment picks up, but because workers leave the state.” Meanwhile Shimer (2007) is more specific, stating that “Blanchard and Katz (1992) find that ten years after an adverse regional shock to employment, higher job creation offsets only about a third of the decrease in employment; *most of the long-run adjustment occurs via worker emigration.*” (emphasis added). Other interpretations appear to rule out firm responses altogether: “Blanchard and Katz (1992) find that in response to local shocks, employment responds in the long run, indicating that labor markets equilibrate through labor mobility” (Kothari, Saporta-Eksten, and Yu, 2013, footnote 2). As we have demonstrated through the examples depicted in Figures 2 and 3, one cannot make any statement regarding the size of the labor demand response based on eye-balling the employment response.

The problem posed by persistent shocks potentially extends to VARs beyond the present application. While SVARs are rightly lauded for their flexibility (see, e.g., Sims, 1980), in the present application the SVAR is subjected to a rather stringent interpretation (especially in comparison to other applications of SVARs in the broader macro literature). Applications in which the long-run response of an endogenous variable is interpreted *relative* to the size of the initial shock will be hampered if that shock is persistent.

2.2 Accounting for Shock Persistency

In this section we present our methodology for decomposing regional recoveries into labor demand and labor migration channels from empirical VECMs. Recall that $\mathbf{\Lambda} \equiv (\lambda^D, \lambda^S)'$ contains the endogenous response parameters of interest. For example, if we are interested in tracing out the

effect of the endogenous labor demand response on employment and the effect of the endogenous labor supply response on population after a certain kind of labor market shock from s periods ago, we have

$$\frac{\partial Y_{i,t}}{\partial \varepsilon'_{i,t-s}} = \mathbf{D}^s \mathbf{\Gamma}, \quad \mathbf{D} \equiv \mathbf{I} + \mathbf{\Gamma} \mathbf{\Lambda} \mathbf{\delta}' \quad (5)$$

where $\mathbf{\Gamma}$ and $\mathbf{\delta}$ are defined as above in section 2, and recall that $Y_{i,t} = (e_{i,t}, p_{i,t})'$. Therefore, in order to decompose a regional recovery into labor demand and supply responses, we require consistent estimates of $\mathbf{\Lambda}$ and $\mathbf{\Gamma}$ from our empirical model (recall that $\mathbf{\delta}$ is a known cointegrating vector). Expressing the structural VECM (4) in reduced-form, we have

$$\Delta Y_{i,t} = \beta Z_{i,t-1} + u_{i,t}, \quad (6)$$

where $\beta \equiv \mathbf{\Gamma} \mathbf{\Lambda}$, $u_{i,t} \equiv \mathbf{\Gamma} \varepsilon_{i,t}$, and $Z_{i,t} \equiv \mathbf{\delta}' Y_{i,t}$ denotes the equilibrium errors ($Z_{i,t} = e_{i,t} - p_{i,t}$ in the case of (4)). In order to recover β from the reduced-form regression (6) it is necessary that the error term is uncorrelated with the regressor vector. (It is also necessary that we impose some form of identification assumption in order to recover $\mathbf{\Gamma}$ from the covariance estimator of the residuals. We address this below in section 2.4.) This condition is obviously violated if the structural shocks $\varepsilon_{i,t}$ are not uncorrelated over time.

However, we can recover a consistent estimator of β if we include lags of the dependent variable in the empirical regression. For simplicity, suppose that the reduced-form shocks $u_{i,t}$ follow a VAR(1) dependence structure:

$$u_{i,t} = \mathbf{B} u_{i,t-1} + v_{i,t}.$$

Substituting the structure into (6) we obtain

$$\Delta Y_{i,t} = \beta Z_{i,t-1} - (\mathbf{B}\beta) Z_{i,t-2} + \mathbf{B} \Delta Y_{i,t-1} + v_{i,t}$$

The model is over-identified and regressors are collinear, since $Z_{i,t-2}$ is a linear combination of

$Z_{i,t-1}$ and $\Delta Y_{i,t-1}$. This means we can omit $Z_{i,t-2}$ and recover β and $\mathbf{B}$ by estimating

$$\Delta Y_{i,t} = \beta Z_{i,t-1} + \mathbf{B} \Delta Y_{i,t-1} + v_{i,t} \quad (7)$$

The parameters which embody firm and worker responses β are separate from the serial dependence parameters $\mathbf{B}$. To generalize the model to permit richer dependence structures in the structural innovations we can include additional lags:

$$\Delta Y_{i,t} = \beta Z_{i,t-1} + \sum_{s=1}^p \mathbf{B}_s \Delta Y_{i,t-s} + v_{i,t}, \quad v_{i,t} = \mathbf{\Gamma} \varepsilon_{i,t} \quad (8)$$

The reduced-form parameters of the model (β and $\mathbf{B}_s$) can be estimated consistently by standard methods (e.g. OLS provided n and T are large). However, the standard impulse response from the fitted VECM is based on the empirical estimates of both β and $\{\mathbf{B}_s\}_{s=1}^p$, meaning that these responses will embody both the endogenous labor demand and supply responses of interest (captured in $\beta \equiv \mathbf{\Gamma} \mathbf{\Lambda}$) as well as the serial dependence in the structural shocks (embodied in the lag matrices $\{\mathbf{B}_s\}_{s=1}^p$). The latter are not lent an economic interpretation in the standard BK model.

Our solution is to model the response of the endogenous variables to a “one-off” shock: i.e. a shock that exhibits no persistence. We do this by computing impulse responses from a shock that exhibits no serial dependence. To do so, when computing IRFs we “turn off” the lagged matrices in the fitted model by setting them to zero. This will give us back the impulse responses described in (5) and that resemble Figure 1 (assuming that the estimates of the elements of β are correctly signed).

Yet another approach would be to incorporate mechanisms into the economic model that predict the observed amplification of the initial labor demand shock that we observe in the employment response. The methodology we propose is rather accommodating and can capture these model extensions provided that the posited mechanisms can be captured by the basic SVAR framework. For example, if capital adjustment costs imply that changes in labor demand depend (in part) on employment growth in the previous period, then we would end up with an economic model that is

nested by (7). In effect, under these theories we are simply applying an economic interpretation to the additional parameters $\{\mathbf{B}_s\}_{s=1}^p$ included in the model.

2.3 Empirical Specification

In our main empirical specification we follow Blanchard and Katz (1992) and extend the economic model in several ways. First, as discussed at length above, we permit lags of the endogenous variables in order to account for potential persistence in the shocks of the economic model. Second, we model three labor market variables (all in logs): Employment, working age population, and labor force (respectively denoted e , l , and p), and we split out the broad employment rate $e - p$ into the “employment rate” $e - l$ and the “participation rate” $l - p$. Note the employment rate (which we model in natural logs) is equal to one minus the unemployment rate, i.e. $\ln\left(\frac{E}{L}\right) = \ln\left(1 - \frac{L-E}{L}\right)$. Both the employment rate $e - l$ and the participation rate $l - p$ will be included as equilibrium errors in the empirical VECM (after conditioning on cross section and time period fixed effects in the empirical specification). Thus we have $Y_{i,t} = (e_{i,t}, p_{i,t}, l_{i,t})'$ and $Z_{i,t} = (e_{i,t} - l_{i,t}, l_{i,t} - p_{i,t})'$.

We differ from Blanchard and Katz (1992) slightly in terms of the specification of our empirical model. While our main empirical specification in the reduced-form VECM given in (8), Blanchard and Katz (1992) adopt a VAR which models employment growth Δe , the unemployment rate $l - e$ and the participation rate $l - p$. The drawback of this approach is that it is difficult to separate the mobility parameters $\mathbf{\Lambda}$ from the serial dependence parameters $\{\mathbf{B}_s\}_{s=1}^p$ in the reduced-form model (i.e., the reduced form parameters are a linear combination of $\mathbf{\Lambda}$ and $\{\mathbf{B}_s\}_{s=1}^p$.) A reduced-form VECM in $\Delta Y_{i,t} = (\Delta e_{i,t}, \Delta p_{i,t}, \Delta l_{i,t})'$ is also a more general model than a reduced-form VAR in $(\Delta e_{i,t}, l_{i,t} - e_{i,t}, l_{i,t} - p_{i,t})$, since there is a mapping of the VAR parameters back to the VECM parameters, but the reverse is not true (see pp. 51-52 of Mark, 2004, for an example).

2.4 Identification

The reduced-form innovations $v_{i,t}$ embody the underlying structural innovations $\varepsilon_{i,t}$ through the mapping $\mathbf{\Gamma}$:

$$v_{i,t} = \mathbf{\Gamma}' \varepsilon_{i,t}, \quad \mathbf{E}(\varepsilon_{i,t} \varepsilon_{i,t}') = \mathbf{\Sigma},$$

where Σ is a diagonal matrix with the variance of the structural shocks along the principal diagonal. Let $\varepsilon_{i,t} = (\varepsilon_{i,t}^d, \varepsilon_{i,t}^{s'})'$. We use two different strategies to recover the labor demand shocks $\varepsilon_{i,t}^d$ from the reduced-form innovations $v_{i,t}$.

The first strategy is to use Bartik's (1991) shift-share employment variables as the labor demand shocks. This approach is often used in empirical models of regional adjustment to identify labor demand shocks (see, amongst others, Saks, 2008; Zabel, 2012). Under this approach the reduced-form errors are assumed to satisfy $v_{i,t} = \eta \varepsilon_{i,t}^{d, \text{BT}} + \xi' \varepsilon_{i,t}^s$, where $\varepsilon_{i,t}^{d, \text{BT}}$ is the observable Bartik shift-share, and $\mathbf{\Gamma}' = [\eta : \xi']$. Thus η can be recovered by including $\varepsilon_{i,t}^{d, \text{BT}}$ in the empirical VECM. Details on how these variables are constructed can be found in the Appendix.

The second method is recursive identification. This approach is also often used in empirical models of regional adjustment to identify labor demand shocks (see, amongst others, Blanchard and Katz, 1992; Decressin and Fatás, 1995). Under this approach Γ is restricted to being an upper-triangular matrix. Because employment is the first element in the $k \times 1$ vector of endogenous variables $Y_{i,t}$, this restriction implies that employment levels are only contemporaneously affected by labor demand shocks. (Equivalently, all other structural shocks, such as labor supply shocks, are precluded from affecting employment levels.) This corresponds to inelastic labor demand in the short run. Under this assumption we can recover Γ using a Cholesky decomposition of the covariance matrix of the residuals obtained from the reduced-form VECM. Recursive identification (or inelastic short run labor demand) necessarily over-estimates the labor supply response if short-run labor demand is in reality elastic, and so provides an upper bound for relative worker mobility. This might be a problem were we to find greater labor supply responses than previously thought or if recursive and Bartik identification methods produced markedly different answers, but as will be shown below, neither of these is the case.

Another reason to explore recursive identification is that the VECM with recursive identification nests both the empirical VAR model used by Blanchard and Katz (1992) and the regressions used by Bartik (1993) to analyze the incidence of employment shocks on regional unemployment and participation rates. This is because both Blanchard and Katz (1992) and Bartik (1993) identify labor demand shocks assuming that short-run changes in employment exclusively reflect innova-

tions to labor demand (and not labor supply), which is equivalent to the recursive identification method described above. That the structural VECM nests the SVAR under these conditions is a straightforward consequence of the reduced-form VECM nesting the reduced-form VAR, as we have already discussed above. Meanwhile Bartik (1991, 1993) estimates a model of the general form of

$$\Delta(l_{i,t} - e_{i,t}) = a_0 \Delta e_{i,t} + \sum_{s=1}^p a_s \Delta e_{i,t-s} + v_{i,t},$$

where $\Delta e_{i,t}$ on the right hand side is used to approximate a labor demand shock. After some simplifications we can establish that $a_s = b_{21}^{(s)} - b_{11}^{(s)}$, where $b_{ij}^{(s)}$ is the (i, j) element of $\mathbf{\Gamma}^{-1} \mathbf{B}_s$, while $a_0 = \gamma_{21}$, where γ_{ij} is the (i, j) element of $\mathbf{\Gamma}$. In this way there is a mapping from the Bartik model to the structural VECM, but not vice versa.

3 Data and Descriptive Statistics

In our main empirical specification we fit the model to a MSA data set consisting of annual employment, labor force and working age population that spans 1990-2012. We obtain the labor market variables at the county level and aggregate up to MSAs based on 2009 MSA definitions. The MSA definitions we use therefore remain consistent over time. We obtain annual wage and salary employment by county from the BEA. Labor force data are constructed as the sum of employment and unemployed. Unemployment numbers are obtained from the Bureau of Labor Statistics (BLS) as part of the “Local Area Unemployment Statistics” data set (LAUS), and are available for counties beginning in 1990. The data represent unemployment by place of residence (not by place of work).⁵ We use annual averages as reported by the BLS.

⁵Although the vast majority of people live within MSAs in the US, we are admittedly omitting rural residents from the sample. However, Partridge *et al.* (2008) find that adjustments in rural labor markets tend to be much more protracted than adjustments in MSA labor markets, suggesting that metropolitan labor market adjustments should be modelled separately from rural labor markets. In addition, some data we use are also only available for MSAs (such as house prices), and the extension of these data to a geographically exhaustive set of labor market definitions such as commuting zones is not straightforward. While our results may nonetheless be confounded by rural residents commuting in to an MSA because they will be included in our employment data but excluded from our unemployment and population data, an unreported robustness check suggests that these commuters are not substantially affecting our results: When we estimate the model with the BLS LAUS place-of-residence employment data, so that these commuters are omitted from the sample, the results are quantitatively similar to those estimated using BEA data.

We digress briefly to discuss why we use these particular data for analyzing regional adjustments at the MSA level. Our objective when selecting a particular data set is to obtain both broad coverage (across time and space) and to minimize potential measurement error in the data. The binding constraint on coverage is MSA-level labor force data. The LAUS data are available starting in 1990 and ending in 2012 (at the time of writing). The only other potential data source with similar coverage is the Geographic Profile of Employment and Unemployment (GP), which is also published by the BLS division of Local Area Unemployment Statistics. Unfortunately, the geography of these data spans neither all MSAs (it only spans a small set of populous MSAs which grows over time) nor all counties. The BLS releases these data annually and does not update past data based on changes in MSA definitions, so that there are serious concerns about developing consistent time series. The data are not available at the county level for good reason—the estimates are based solely on the current population survey (CPS), which consists of too small a sample of individuals to produce quality county-level estimates.

In addition to the BEA there are many more sources for annual employment figures at the county level (including the LAUS; the Quarterly Census of Employment and Wages, or QCEW; and the Employment and Earnings, or E&E, which is an older publication of the BLS division of Local Area Unemployment Statistics). We use the BEA employment data for three reasons. First, they are almost certainly the most accurate source of county-level employment data available. The BEA incorporates several data sources (including the QCEW) to ensure complete coverage and makes adjustments for systematic misreporting. This means that BEA data are more comprehensive than QCEW data, which only cover employees that are part of state unemployment insurance (UI) programs. In contrast, the LAUS combines data from different sources within a state-space model to produce county-level estimates of employment and unemployment⁶, meaning that there is considerable measurement error in these data. E&E data are based on a survey (not a census such as QCEW) and are not available at the county level. The E&E data were originally used by Blanchard and Katz (1992). Rowthorn and Glyn (2006) argue that measurement error in the Blanchard and Katz (1992) data set leads to an understatement of the persistency in state level

⁶For details see <http://www.bls.gov/lau/laumthd.htm>

employment rates, and thus adoption of BEA data is likely to represent a substantial improvement over previous work and over other data source options. Second, the BEA data are defined on a place-of-work basis, which is consistent with our model, while LAUS data are defined on a place-of-residence basis. Regardless, our results do not substantially change when the LAUS employment data are used in place of the BEA employment data (these results are unreported, but are available upon request). Third, we estimate a more parsimonious version of the model (which omits the labor force from the specification) over a longer time-span (1970-2012) in order to examine whether regional adjustment channels have changed over time. Any differences in the results of the two data sets cannot be attributed to differences in the underlying employment data source.

It is important to note then that our measure of the labor force (the sum of employment and unemployed) likely has more measurement error than employment since we obtain county-level unemployment data from the LAUS. We use the LAUS for unemployment as it is the only county-level data set with sufficiently broad coverage. As noted above, however, the only other source for subnational unemployment statistics known to the authors is the GP, which was used by Blanchard and Katz (1992). This data set is based exclusively on the rather sparse CPS and does not contain county-level information. Thus, while measurement error is unavoidable given a lack of usable alternative measures of unemployment, the data used in this paper are undoubtedly less noisy than those to which Blanchard and Katz (1992) had access.

We infer net migration responses by looking at changes in working-age population. We use the working-age population (as opposed to total population) as it means our results are less likely to be contaminated by long term demographic shifts (such as the retirement of the baby boomers). Working-age population by county data are Census Bureau mid-year age 15-64 postcensal and intercensal population estimates. We do however face a trade-off in adopting working age population instead of total population because the former is measured with less accuracy than the latter. Unreported results based on total population are however nearly identical to those based on working-age population, suggesting that this particular measurement issue has little effect on our findings.

The Census Population Estimates Program uses the IRS income tax return data (among other sources) to annually calculate internal migration data for postcensal populations at the state, county,

and county equivalent level. Census therefore has access to the age of the filer when computing migration rates. Furthermore, the Census collects migration information from a number of different sources as they come available, such as the American Community Survey (ACS), and used different sources prior to the IRS migration data becoming available. Thus, these data incorporate the best available migration information at any given time, and extend over a long time horizon. The best alternative would be to use IRS tax-return-based migration data directly, which we include in a robustness check (it does not make a notable difference). However, IRS data come with many additional caveats relating to measurement error, and hence they are not used in our main empirical specification. For example, because they are based only on households that file income taxes, the low income households are likely to be under-represented. Divorces, marriages and adolescents leaving the household are also missed.

We also consider extensions of the basic model that include nominal house prices, wages, housing supply elasticities, direct measures of migration, and the number of establishments. All data used in our main empirical specification and extensions are summarized in table 1. Additional details on the data sources and construction can be found in the Appendix.

Table 1: Data and sources

Variable	Data source ¹	Institution ¹	Geography	First year	No. MSAs ²
Employment	Local area empl.	BEA	County	1969	364
Unemployment	LAUS	BLS	County	1990	363
Population (15-64 years)	Intercensal pop.	Census	County	1970	364
Wages	Local area wages	BEA	County	1969	364
House prices	House price index	FHFA	MSA	various	326
Housing supply elast.	WRLURI	Wharton	MSA	n/a	257
Migration	SOI Tax Stats	Census	County	1990	364
Establishments	CBP	Census	County	1974	364
Bartik variable	QCEW	BLS	County	1990	364

¹ Abbreviations: BEA: Bureau of Economic Analysis. LAUS: Local area unemployment statistics. BLS: Bureau of Labor Statistics. FHFA: Federal Housing Finance Agency. WRLURI: Wharton Residential Land Use Regulatory Index. SOI: Statistics of Income. CBP: County Business Patterns. Census refers to the US Census Bureau. QCEW: Quarterly census of employment and wages

² When data from different sources are combined, only the MSAs contained in the intersection of the sources are retained. The number reported in this column represents the largest number of MSAs that are actually used when the variable is included in the model.

3.1 Descriptive Statistics

Table 2 provides some basic descriptive statistics of the data. Employment growth was barely positive on average across all MSAs over the 1990-2012 period. Note that this includes the recent fall in US employment resultant from the 2007-2009 recession. Both the labor force and working age population grew at just over 1% per year on average. The average (log) participation rate of -0.288 corresponds to a mean participation rate of 75.0% in levels (note we are using working age population in constructing this rate), while the average (log) employment rate of -0.070 corresponds to an employment rate of 93.2% in levels (or equivalently an unemployment rate of 6.8%). “std dev” denotes the average standard deviation across each time series in the panel. In this column, we can see that employment growth is about 40% more volatile than labor force growth, and is about two and a half times more volatile than is working age population growth. Average correlation (“avr. corr.”) gives an indication of the cross sectional dependence in the panels, and hence the potential affect of aggregate shocks on each of the variables. This number is generated by computing correlations between cross-sections and averaging their absolute values across pairs of locations. Employment growth exhibits a larger cross sectional correlation than labor and population growth. Similarly, the employment rate exhibits much more correlation between MSAs than the participation rate.

Table 2: Descriptive Statistics: MSA labor market data 1990-2012

MSA data set	mean	std. dev.	mean abs. corr.	pooled AR(1)
employment growth (Δe)	0.001	0.021	0.510	0.320
working age population growth (Δp)	0.012	0.008	0.333	0.423
labor force growth (Δl)	0.010	0.015	0.273	0.266
employment rate ($e - l$)	-0.070	0.022	0.728	0.822
participation rate ($l - p$)	-0.288	0.030	0.481	0.880

“std. dev” and “mean abs. corr.” denote the mean standard deviation across each time series in the panel and the mean absolute correlation coefficient between the time series in the panel. “AR1” denotes the AR(1) coefficient fitted by ordinary least squares, with period and cross section fixed effects included.

We also report OLS estimates of the pooled AR(1) coefficients to give an indication of the serial

dependence in each panel. These are estimated based on a simple AR(1) model with cross section and period fixed effects. The employment and participation rates exhibit sufficient persistence to warrant more formal testing for unit roots. In this regard we require a unit root test that is appropriate for the features of our data set. First, as demonstrated in Table 2, there is significant cross sectional correlation between the time series of the panel, and cross sectional dependence can induce a size distortion in conventional panel data unit root tests that require cross sectional independence (such as the well-known Levin, Lin and Chu (2002) and Im, Pesaran and Shin (1999) tests; see chapter 12 of Baltagi, 2013, for a general discussion). Second, the time series dimension of the panel is sufficiently limited to make small-sample bias a concern, yet sufficiently large to invalidate methods that are based on so-called “fixed T” asymptotics ($n \rightarrow \infty$ only). Instead, we require methods that permit both n and T to be large. We therefore use the panel unit root tests proposed by Hosseinkouchack and Westurlund (2016), which make a simple adjustment to the popular CIPS and CADF tests proposed by Pesaran (2007) to ensure that the test statistics have standard limiting distributions. This family of tests account for cross sectional dependence via a factor structure in the error term, and exhibit the correct size in tests of the null of a unit root in simulated panel samples of similar size to our data set (Hosseinkouchack and Westurlund, 2016).

The modified CIPS test statistics are 15.74 and 5.62 for the employment and participation rates.⁷ These test statistics have a standard normal limiting distribution under the null of a unit root, and thus we can reject the null in both employment and participation rates at conventional significance levels. Thus, although MSA participation (and to a lesser extent employment) rates are quite persistent, they do appear to be stationary after conditioning on period and cross-section fixed effects. This contrasts with regional unemployment and participation rates at the state level, for which there is little evidence in favor of stationarity, particularly when period effects are omitted (see Table 2 in Partridge *et al.*, 2015). A variety of explanations could account for the differences in the time series properties of unemployment rates at different levels of regional aggregation. For example, aggregation of individual time series with different serial dependence dynamics could result in greater persistence in the aggregate. It is important to note, however, that measurement error

⁷Lag orders for the augmented Dickey-Fuller regressions are selected using the Schwarz criterion.

could be biasing the findings towards rejecting the null of a unit root (see, e.g., Rowthorn and Glyn, 2006). As we stress above, we carefully selected our data sources in our to minimize the potential measurement error in our labor market variables.

4 Empirical Results

In this section we present our empirical findings. First we consider the standard BK empirical model applied to MSA labor market variables (employment, labor force, and working age population) over the time period from 1990 to 2012. We demonstrate that the identified labor demand shocks exhibit serial dependence. We cannot, therefore, determine the contribution of the labor demand response to the recovery by looking at the employment response. Then, in subsection 4.2 we present the imputed impulse responses with the serial dependence parameters “switched off”. Subsection 5 provides extensions and robustness tests of this baseline model and data set, including a longer MSA data set and model specifications that include wages and house prices, regional housing supply elasticities, and direct measures of migration.

In all VECMs the lag order p is selected using the Lee and Phillips (2015) integrated likelihood information criterion (ILIC), which is a generalization of the better-known Schwarz criterion to a panel data setting. Because we do not know of a multivariate version of the ILIC, we compute the criterion for each equation within the VECM and then set p to be the maximum lag order selected by the ILIC across all equations. Unless otherwise stated, this results in two lags in most regressions.

4.1 Regional Population and Employment Responses to Serially Dependent Labor Demand Shocks

In this subsection we present the response of employment, labor force and population to labor demand shocks without regard for potential serial dependence in the shocks.

To account for long run differences in regional unemployment and participation rates, we augment the model (8) with (a vector of) cross sectional fixed effects γ_i . Because relative regional

employment and participation rates are the relevant state variables, we also include (a vector of) time period fixed effects (or common time effects) d_t . Thus the empirical model is

$$\Delta Y_{i,t} = \beta Z_{i,t-1} + \sum_{s=1}^p \mathbf{B}_s \Delta Y_{i,t-s} + \gamma_i + d_t + \eta \varepsilon_{i,t}^{d, \text{BT}} + v_{i,t} \quad (9)$$

Impulse responses to labor demand shocks are obtained via the following steps:

1. Estimate (9) by OLS.
2. Compute the impulse response functions for the labor demand shocks based on the fitted model. Labor demand shocks are normalized to generate a -1% fall in employment in the first period.
3. Construct confidence intervals for the IRFs based on Efron’s (1981) bootstrapped standard errors. See section A.2 in the Appendix for details.

Figure 4 exhibits the estimated impulse response functions for employment, labor force and working-age population up to twenty years after the initial shock.⁸ These IRFs are also tabulated in Table 3 below at short term (1 year), medium term (5 and 10 year) and long run (20 year) horizons.

The employment response exhibits the familiar inverse “hump” shape documented in other empirical models of regional adjustment. Employment continues to fall for an additional two periods after the initial -1% shock, but eventually recovers to about -0.64% after twenty years. This long run effect of the labor demand shock on employment is notably less than that found in the comparable literature: Using MSA data, Saks (2008) and Zabel (2012) find that a labor demand shock that initially decreases employment by 1% results in a permanent 1% reduction in employment over the long run; BK found a similar result using state level data. Nonetheless, Figure 4 suggests that a large amount of household migration is generated by the shock. However, because we cannot

⁸Note that the cumulated impulse responses will be particularly sensitive to specification error at longer horizons. Therefore it may be somewhat heroic to extrapolate our results out to the twenty-year horizon. We do however attempt to mitigate this problem by using formal model selection methods designed for panel data models. In addition, the results remain largely unchanged when we estimate the model using larger lag orders.

observe the increase in local employment due to the labor demand response, we cannot compare the relative magnitude of the labor supply response (out-migration) to the labor demand response.

Figure 5 exhibits the impulse responses for the employment and participation rates. While unemployment (one minus the employment rate) recovers relatively quickly (within ten years), the participation rate takes in excess of twenty years to recover.

4.2 Regional Population and Employment Responses to Serially Uncorrelated Labor Demand Shocks

In this subsection we present our baseline results the response of employment, labor force and population to a one-period labor demand shock. Impulse responses to labor demand shocks are obtained via the following steps:

1. Estimate (9) by OLS.
2. Compute the impulse response functions for the labor demand shocks based on the fitted model, but set $\hat{\mathbf{B}}_s = \mathbf{0}$ for all $s = 1, \dots, p$. Labor demand shocks are normalized to generate a -1% fall in employment in the first period.
3. Construct confidence intervals for the IRFs based on Efron's (1981) bootstrapped standard errors. See section A.2 in the Appendix for details.

The impulse responses of employment, population and labor force are shown in Figure 6. These IRFs are also tabulated in Table 3 below at short term (1 year), medium term (5 and 10 year) and long run (20 year) horizons.

Employment falls by 1% in the first period. After this initial shock, population continues to decrease and employment increases, respectively reflecting the effects of the endogenous labor supply and demand responses to the negative labor demand shock. Importantly, and in contrast to Figure 4 above, employment does not continue to decrease after the initial fall. The recovery begins immediately, in line with the BK model. After twenty years employment has decreased by

Figure 4: Responses to a serially correlated -1% labor demand shock using Bartik shift-shares. Error bands represent 90% confidence intervals.

Figure 5: Responses to a serially correlated -1% labor demand shock using Bartik shift-shares. Error bands represent 90% confidence intervals.

0.37%, meaning that employment in fact recovered by 0.63% ($= 1\% - 0.37\%$) in the years following the initial 1% fall in employment. Because the initial shock exhibits no serial dependence, this subsequent increase in employment of 0.63% is due to the labor demand response. Meanwhile, population has declined by 0.24% twenty years after the shock. Over the long run, the increase in employment due to the labor demand response is about two and a half times larger than the out-migration due to the labor supply response ($2.5 \approx 0.63 \div 0.24$). Firms therefore contribute more to the recovery in the participation and unemployment rates than workers.

The impulse responses of the participation and employment rates are shown in Figure 7. In the first period the IRFs reflect the initial incidence of the shock on the participation and employment rate margins. From the second period onwards these IRFs exclusively reflect the effect of labor supply and demand responses to the shock, thereby telling us about the extent of the recovery in the employment and participation rates in each period subsequent to the shock.

About two-thirds of the labor demand shock is initially absorbed through a decrease in the participation rate, and just less than a quarter is absorbed through a decrease in the employment rate (equivalently, an increase in the unemployment rate). The remainder of the shock is accounted for through the out-migration channel in the first period, as shown in the second panel of Figure 6. The effect of the labor demand shock on the employment and participation rates is long-lasting. The participation rate in particular remains significantly below its pre-shock level twenty years after the original shock.

Our method of filtering out shock persistence bears some similarities to an earlier exercise performed by Bartik (1993) in reconsidering Blanchard and Katz (1992) and Bartik (1991). Bartik (1993) models a “fixed” shock to employment by not directly allowing the shock to affect subsequent employment growth in the Blanchard and Katz (1992) VAR. While this method precludes shock persistence, it will also preclude the endogenous labor demand response we are interested in studying, since employment levels are not permitted to adjust after the shock. Bartik (1993) also demonstrates that the Blanchard and Katz (1992) findings are not robust to various specification issues, and in particular finds that these “fixed” employment shocks can have very persistent effects on the participation rate when more lags are included in their model. Our results are far more

Figure 6: Responses to a serially uncorrelated -1% labor demand shock using Bartik shift shares. Error bands represent 90% confidence intervals.

Figure 7: Responses to a serially uncorrelated -1% labor demand shock using Bartik shift shares. Error bands represent 90% confidence intervals.

Table 3: Estimated Impulse Response Functions

period (years)	Baseline model (no adjustment for serial dependence in structural shocks)					
	Participation rates	Employment rates	Employment	Working-age population	Labor force	
1	-0.73 (-0.61, -0.85)	-0.20 (-0.09, -0.31)	-1.00 (-1.00, -1.00)	-0.07 (-0.01, -0.13)	-0.80 (-0.69, -0.91)	
5	-0.53 (-0.43, -0.64)	-0.13 (-0.06, -0.20)	-1.15 (-0.96, -1.35)	-0.49 (-0.38, -0.60)	-1.02 (-0.88, -1.17)	
10	-0.28 (-0.19, -0.37)	-0.06 (0.01, -0.12)	-0.87 (-0.65, -1.10)	-0.54 (-0.40, -0.68)	-0.81 (-0.63, -0.99)	
20	-0.08 (-0.04, -0.13)	-0.02 (0.00, -0.03)	-0.64 (-0.45, -0.82)	-0.54 (-0.38, -0.70)	-0.62 (-0.45, -0.80)	
period (years)	Baseline model (with serial dependence in structural shocks filtered out)					
	Participation rates	Employment rates	Employment	Working-age population	Labor force	
1	-0.73 (-0.61, -0.85)	-0.20 (-0.08, -0.32)	-1.00 (-1.00, -1.00)	-0.07 (-0.01, -0.13)	-0.80 (-0.68, -0.92)	
5	-0.50 (-0.44, -0.56)	-0.12 (-0.08, -0.16)	-0.75 (-0.68, -0.83)	-0.13 (-0.08, -0.19)	-0.63 (-0.57, -0.70)	
10	-0.30 (-0.24, -0.37)	-0.07 (-0.04, -0.10)	-0.56 (-0.45, -0.66)	-0.19 (-0.12, -0.25)	-0.49 (-0.40, -0.57)	
20	-0.11 (-0.07, -0.15)	-0.02 (-0.01, -0.04)	-0.37 (-0.27, -0.47)	-0.24 (-0.16, -0.31)	-0.35 (-0.25, -0.44)	
period (years)	Model including establishments data (with serial dependence in structural shocks filtered out)					
	Participation rates	Employment rates	Employment	Working-age population	Labor force	Establishments
1	-0.71 (-0.59, -0.83)	-0.22 (-0.11, -0.33)	-1.00 (-1.00, -1.00)	-0.07 (-0.02, -0.13)	-0.78 (-0.67, -0.89)	-0.14 (0.00, -0.27)
5	-0.50 (-0.44, -0.56)	-0.12 (-0.08, -0.16)	-0.75 (-0.68, -0.82)	-0.13 (-0.08, -0.18)	-0.63 (-0.56, -0.69)	-0.12 (0.01, -0.25)
10	-0.31 (-0.24, -0.37)	-0.07 (-0.04, -0.10)	-0.55 (-0.45, -0.66)	-0.18 (-0.12, -0.24)	-0.48 (-0.40, -0.57)	-0.11 (0.03, -0.25)
20	-0.11 (-0.07, -0.16)	-0.02 (-0.01, -0.04)	-0.36 (-0.26, -0.47)	-0.23 (-0.15, -0.30)	-0.34 (-0.25, -0.43)	-0.09 (0.05, -0.24)

Impulse responses to a 1% decrease in employment in period one. Bootstrapped 90% confidence intervals in parentheses

robust to model specification concerns in that the responses do not change substantially when we alter the lags in the VECM (in the interests of brevity we have omitted these results from the manuscript). However, much like Bartik (1991, 1993), we find that labor demand shocks can have large and long-lasting effects on local residents. For example, twenty years on from the negative 1% shock, the participation rate is 0.11 % lower than its pre-shock level, suggesting that the shock has a large and persistent effect on the participation of local residents in the labor force.

4.3 Decomposing the Labor Demand Response

Our main result highlights the importance of the endogenous labor demand response to the recovery. However, we cannot tell whether the increase in employment is being driven by a net increase in firm formation at the local level or a net increase in employment at existing firms. We explore this issue by reestimating the model with the number of establishments by MSA (denoted as $m_{i,t}$) included in the regression. County-level establishment data is obtained from the Census Bureau and aggregated up to the MSA level based on 2009 MSA definitions.⁹

The empirical model is given by (9), where now $Y_{i,t} = (e_{i,t}, p_{i,t}, l_{i,t}, m_{i,t})'$ but $Z_{i,t}$ remains $Z_{i,t} = (e_{i,t} - l_{i,t}, l_{i,t} - p_{i,t})'$. Note that the dimensions of the reduced-form model parameters β , $\mathbf{B}_s$, γ_i and d_t are re-defined to be conformable.

The impulse responses of employment, population, labor force and establishments are shown in Figure 8. These IRFs are also tabulated above in Table 3. The number of establishments falls with the initial shock, suggesting that the initial downturn coincides with a net increase in firm deaths. The number of establishments subsequently recovers, but only slightly, suggesting that the bulk of the labor demand response is attributable to extant firms expanding their payrolls. Note however that the 90 percent confidence interval for the establishments response includes zero at all horizons, so the precision of the estimates is rather limited. Note, however, that the establishments response obtained under recursive identification of the labor demand shock is estimated with much great precision (see Figure 15 in the Appendix). As shown in Table 3, the recoveries in the participation

⁹We do not explore whether the net increase in firm formation is driven by an increase in the birth rate of new firms or a fall in the death rate of existing firms. Similarly, we cannot tell if the net increase in net employment is driven by a decline in the rate of layoffs or an increase in the the rate of hires

and unemployment rates are largely unaffected by the inclusion of establishments in the VECM.

5 Robustness

In this subsection we consider extensions of the baseline model and methodology given in the preceding subsection. In all cases, we find population and employment responses that are very similar to those estimated using our baseline model and data set. First, we consider alternative empirical VECMs that accommodate data related to housing, wages and inter-MSA migration. Second, we consider an alternate method for identifying labor demand shocks based on recursive identification. Third, we estimate the baseline model using instrumental variables instead of OLS. All relevant figures are contained within the Appendix.

5.1 Alternative Empirical Specifications

We consider several extensions of the basic VECM. First, we consider how migration frictions may affect our results. We include two additional price measures in the model: Wages and house prices. We also use an index of urban land regulation to measure how migration responses may vary across cities with different housing supply elasticities. Second, we consider using direct measures of MSA migration to estimate the migration response. See Table 1 above or the Appendix for additional details on these data.

5.1.1 Housing Price and Wages Specification

House prices and wages are often included in empirical models of regional adjustment in order to include a broader range of economic variables that are likely to affect the migration decision (Saks, 2008; Zabel, 2012). In this subsection we explore whether the addition of these variables to our baseline model affects our results in any significant way.

Let $w_{i,t}$ and $h_{i,t}$ denote the log wage and housing price at time t in location i . The empirical model is given by (9), but we use $Y_{i,t} = (e_{i,t}, p_{i,t}, l_{i,t}, h_{i,t}, w_{i,t})'$. $Z_{i,t}$ remains defined as $Z_{i,t} = (e_{i,t} - l_{i,t}, l_{i,t} - p_{i,t})'$. The dimensions of the reduced-form model parameters β , $\mathbf{B}_s$, γ_i and d_t are

Figure 8: Responses to a serially uncorrelated -1% labor demand shock using Bartik shift shares. Error bands represent 90% confidence intervals.

re-defined to be conformable. The available data for wages and house prices by MSA are more limited, and thus our sample is restricted to 326 MSAs over 1990 to 2012.

Figure 10 exhibits the responses to a -1% labor demand shock. The response of the labor market variables follow similar paths to the specification without wages and housing (exhibited in Figures 6 and 7 above). The long-run change in population is however slightly smaller in magnitude (population declines by 0.2% after twenty years), and the long-run recovery in employment is slightly larger at 0.67% ($= 1\% - 0.37\%$). Thus the relative size of the labor demand response becomes even larger when house prices and wages are included in the empirical model. On the one hand, this could indicate these prices are important when modelling migration decisions. On the other hand, there may be sample selection biases at work because the number of MSAs in this sub-sample is limited.

5.1.2 Regional Heterogeneity in Housing Supply Elasticity

The elasticity of housing supply may be a crucial factor in the ability of a labor market to accommodate new workers, and recent work by Saks (2008) and Zabel (2012) considers how regional variation in housing supply elasticity affects labor market adjustment. To explore this possibility we split the MSAs into quartiles based on the Wharton Residential Land Use Regulation Index (WRLURI) developed by Gyourko, Saiz and Summers (2008) before fitting the VECM model to each quartile.¹⁰ The WRLURI is used as a proxy for housing supply elasticities. We retain 243 MSAs in the more limited WRLURI sample.

Figures 11 and 12 exhibit the responses to a -1% labor demand shock. Note that each figure represents a different quartile. Figure 12 depicts the fourth quartile, which corresponds to MSAs with the highest levels of housing regulation in the WRLURI and thus the lowest levels of supply elasticity. The housing stock within this quartile should not easily expand (or contract) to accommodate migrants. Indeed, we see that although the population decreases slightly (by about 0.1% in

¹⁰We use the WRLURI because it accounts for several factors that may affect housing elasticity, including regulation restrictions and land constraints. The index is based on surveys from 2005. Admittedly this is towards the end of our sample, and thus the elasticity measure is likely to be endogenous, with cities that grew having more regulation. Due to data constraints, this is however a rather intractable problem unless we are willing to severely reduce our sample size.

the long run), the change in population at all horizons is statistically indistinguishable from zero. Figure 11 depicts the first quartile, which corresponds to MSAs with the highest levels of housing supply elasticity. As we would expect, there is a much larger long run migration response for these MSAs. However, the long run recovery in the employment level after the initial 1% fall still exceeds the long run reduction in population even for these MSAs with the most elastic housing supply. We conclude that housing supply elasticities do affect household migration, but that the basic message of the paper remains unchanged: The labor demand response contributes more to the recovery than out-migration.

5.1.3 Direct Measures of Migration

The IRS provides estimates of county-to-county migration based on the zip-code of tax return filings and the number of exemptions claimed. We use this direct measure of migration to consider how workers respond to regional conditions. Let $NM_{i,t}$ denote the number of net migrants to MSA i between year $t - 1$ and t . The empirical model is given by (9), with (with some abuse of notation) $\Delta Y_{i,t} = (\Delta e_{i,t}, \ln(NM_{i,t} + P_{i,t-1}) - \ln(P_{i,t-1}), \Delta l_{i,t})$ but $Z_{i,t}$ remains as $Z_{i,t} = (e_{i,t} - l_{i,t}, l_{i,t} - p_{i,t})'$. As above, we consider the impulse responses to a -1% labor demand shock, and these are exhibited in Figure 13 located in the Appendix.

On the whole the results are strikingly similar to those based on the working-age population (given in Figures 6 and 7 above). However, the long-run population response is slightly larger, which is likely due to the fact that the net migration number $NM_{i,t}$ reflects all members of a migrating household, not just the members of the household who are of working age. (Recall that the IRS measure of migration is based on the number of tax return exemptions, but our working-age measure of population only includes individuals aged 15-64.)

5.2 Recursive Identification of Labor Market Shocks

In this section we use an alternative identification method for recovering labor market shocks from the reduced form VECM. The empirical model is

$$\Delta Y_{i,t} = \beta Z_{i,t-1} + \sum_{s=1}^p \mathbf{B}_s \Delta Y_{i,t-s} + \gamma_i + d_t + v_{i,t} \quad (10)$$

The structural shocks are recovered from the reduced form residuals through recursive identification (recall that employment growth is the first element of $\Delta Y_{i,t}$). We refer the reader to subsection 2.4 for the details.

Impulse responses to labor demand shocks are obtained via the following steps:

1. Estimate (10) by OLS.
2. Apply a Cholesky factorization to the variance-covariance matrix of the reduced-form residuals in order to identify the labor demand shocks.
3. Compute the impulse response functions for the labor demand shocks based on the fitted model, but set $\hat{\mathbf{B}}_s = \mathbf{0}$ for all $s = 1, \dots, p$. Labor demand shocks are normalized to be -1%.
4. Construct confidence intervals for the IRFs based on Efron's (1981) bootstrapped standard errors. See section A.2 in the Appendix for details.

The impulse responses of employment, population and labor force are shown in Figure 14. The recovery paths in employment, population and labor force are similar to those obtained through Bartik shift-share identification. Over the long run, the recovery in employment after the initial -1% shock is approximately twice as large as the decrease in population, meaning that the labor demand response is much larger than the labor supply response. Note however that the labor migration response is much faster than the labor demand response under recursive identification: Approximately half of the long-run out-migration of workers occurs within the first period, whereas it takes about ten years for half of the long-run increase in employment to be realized. This likely reflects the fact that recursive identification attributes most of the short-run covariation between

the variables to the first shock in the ordered vector of variables (which is the labor demand shock in our case).

We also explore the alternative empirical specifications outlined above with recursive identification.

5.2.1 Housing Price and Wages Specification

Recall that $w_{i,t}$ and $h_{i,t}$ are the log wage and log housing price at time t in location i . The empirical model is given by (10), but we use $Y_{i,t} = (e_{i,t}, p_{i,t}, l_{i,t}, h_{i,t}, w_{i,t})'$, while $Z_{i,t} = (e_{i,t} - l_{i,t}, l_{i,t} - p_{i,t})'$. Note that the dimensions of the reduced-form model parameters β , $\mathbf{B}_s$, γ_i and d_t are re-defined to be conformable. Figure 16 in the Appendix exhibits the responses to a -1% labor demand shock. These responses follow similar paths to the specification without wages and housing (exhibited in Figure 14). However, the long-run decline in population is slightly smaller, and the recovery in employment is slight larger, meaning that the relative size of the labor demand response becomes even larger when house prices and wages are included in the empirical model.

5.2.2 Regional Heterogeneity in Housing Supply Elasticity

Figures 18 and 19 exhibit the responses to a -1% labor demand shock in the 1st (highest elasticity) and 4th (lowest elasticity) WRLURI quartiles, respectively. We can see that the long-run decline in population is larger in the high elasticity quartile (Figure 18), again suggesting that housing supply elasticities affect household migration. Nonetheless, in both quartiles the increase in employment after the initial shock is larger than the decline in population over the long run.

5.2.3 Direct Measures of Migration

The empirical model is given by (10), but with $\Delta Y_{i,t} = (\Delta e_{i,t}, \ln(NM_{i,t} + P_{i,t}) - \ln(P_{i,t-1}), \Delta l_{i,t})$ and $Z_{i,t} = (e_{i,t} - l_{i,t}, l_{i,t} - p_{i,t})'$, and where recall that $NM_{i,t}$ denotes the number of net migrants to MSA i between year $t - 1$ and t . Impulse responses to a -1% labor demand shock obtained from recursive identification are exhibited in Figure 20 located in the Appendix. Results are very similar to those based on the working-age population (given in Figure 14), but the long-run population

response is slightly larger. As discussed above, this is likely due to the fact that the net migration number $NM_{i,t}$ reflects all members of a migrating household, not just the members of the household who are of working age.

5.2.4 Longer MSA panel

Admittedly our baseline results in subsection 4.2 are based on a short time span: 1990-2012. This is due to the restriction on labor force data available at the MSA level. It is possible, for example, that our conclusions differ from studies that analyzed earlier time periods simply because local areas have responded differently to shocks in the most recent two decades than they have to shocks in previous decades.

To rule this out, we consider a more parsimonious specification that omits labor force from the model. We instead take the “broad employment rate” (the sum of the employment and participation rates) as the indicator of local economic conditions. Thus $Y_{i,t} = (e_{i,t}, p_{i,t})'$, and the cointegrating error becomes $Z_{i,t} = e_{i,t} - p_{i,t}$. Our sample of employment and working-age population spans 1970 to 2012.¹¹ The ILIC selects three lags for the VECM.

Figure 17 in the Appendix exhibits the response paths of employment and population. The recovery speeds and shares are very similar to those exhibited for the shorter sample in Figure 14, but the recovery period for the broad employment rate is slightly more protracted. But on the whole these results suggest that firm and worker responses have not changed much since the 1970s. Note that this finding is not inconsistent with declining gross migration rates in the US over the past four decades (See Molloy, Smith and Wozniak, 2011; Kaplan and Schulhofer-Wohl, 2012), since declining migration is accounted for in the BK model through a decrease in the variance of labor market shocks.

¹¹Unfortunately, due to limitations on employment by industry and county, we cannot use Bartik shift-share identification in this longer dataset.

5.3 Instrumental Variables Estimation

OLS bias may be a concern given that the time series dimension of the panel is rather limited (22 annual observations). To investigate whether this bias has a material effect on our results we estimate the baseline VECM specification (9) using the Hayakawa (2009) instrumental variables (IV) estimator. The Hayakawa estimator is optimal in the sense defined by Arellano (2015) for “large n and T ” panels (also see the discussion on p. 170 of Arellano, 2003), and, unlike other IV or limited information maximum likelihood estimators, the estimator does not exhibit bias in the limiting distribution (Hayakawa, 2009).

Figure 21 in the Appendix exhibits the recovery paths when the shocks are identified using the Bartik shift-shares. In comparison to the responses obtained from the OLS estimator (depicted in Figures 6 and 7), we observe a larger migration response and smaller recovery in employment: After twenty years, population declines by 0.33%, while there is a $0.45\% = 1\% - 0.55\%$ recovery in employment. In addition, the recovery is more protracted. Twenty years after the shock, the participation rate remains -0.17% below its initial level. Nonetheless, these results continue to support the notion of a recovery primarily driven by the labor demand response.¹²

Figure 22 in the Appendix exhibits the recovery paths when the shocks are identified using recursive identification. We observe a larger long run migration response, a smaller recovery in employment, and a more protracted recovery compared to the model estimated by OLS (depicted in Figure 14). Nonetheless these results continue to suggest a recovery primarily driven by the labor demand response.

6 What Generates Persistence in Labor Demand Shocks?

In this section we shed light on the underlying causes of shock persistence and amplification by exploring the extent to which local input-output effects are generating the observed amplification in identified labor demand shocks. To do so, we estimate a version of the empirical VECM (9) in

¹²The fitted model implies that population and employment decline by 0.41% in the long run. Thus the effect of the labor demand response on employment is about one and a half times larger than the effect of the labor supply response on population ($1.5 \approx 0.59\% \div 0.41\%$).

which we include industry-level employment growth in the vector of dependent variables. We may uncover the posited input-output effects by examining the response of industry-level employment to a labor demand shock in another sector.

We disaggregate MSA employment data into six constituent industry groups (henceforth “sectors”): Manufacturing and Mining (“MM”); Professional and Other Services (“PS”); Education and Health Services (“EH”); Leisure (“L”); Transport, Trade and Utilities (“TTU”); and Construction (“C”). These six sectors each account for a similar proportion of total employment on average, except for Construction, which is left as a separate category due its volatility. (MM accounts for 19% of employment on average, across all MSAs and years; PS accounts for 23%; EH accounts for 15%; L accounts for 12%; TTU accounts for 24%, and C accounts for 6%.) We refer the reader to the Appendix for further details on the employment data.

We then estimate (9) with $Y_{i,t} = (e_{i,t}^{\text{MM}}, e_{i,t}^{\text{PS}}, e_{i,t}^{\text{EH}}, e_{i,t}^{\text{LH}}, e_{i,t}^{\text{TTU}}, e_{i,t}^{\text{C}}, e_{i,t}, p_{i,t}, l_{i,t})'$, where $e_{i,t}^{\text{MM}}$ denotes log manufacturing and mining employment, for example, and where the cointegrating matrix β is defined such that $Z_{i,t} = \delta' Y_{i,t} = (e_{i,t} - l_{i,t}, l_{i,t} - p_{i,t})$. We include log MSA employment $e_{i,t}$ in the VECM to allow computation of the equilibrium error employment rate $e_{i,t} - l_{i,t}$. The Lee and Phillips (2015) ILIC selects one lag.

We report the responses of the six sectoral employment rates to a labor demand shock to manufacturing and mining identified using a Bartik shift-share variable constructed for the MM sector. We focus on manufacturing and mining under the assumption that a significant proportion of the production is exported, and therefore the Bartik shift share reflects exogenous variation in the labor demand for this sector. While sectors such as financial, information, professional and business, and other services also likely export some production, it is also likely that a significant proportion of production is consumed within the MSA.

Figure 9 exhibits the response to a labor demand shock to manufacturing and mining. Several features of the IRFs indicate input-output effects generate the observed persistency in shocks to aggregate MSA employment. First, the sector-specific MM labor demand shock exhibits little serial dependence: although employment continues to decline in the second period to -1.1%, it begins to recover after the second period, and eventually recovers to -1% in the long run. The inertia in

the shock is therefore much less compared to that based on total MSA employment (see Figure 5 above). Second, we can see that a labor demand shock to the MM sector tends to generate an immediate decrease in all other sectors except Education and Health. More importantly, the shock continues to generate reductions in these other sectors in the second, third and fourth years after the shock. For example, construction continues to decrease over the second through fourth years after the shock, and only begins to recover in the fifth year. Financial services and TTU also exhibit sustained decreases in employment in the second, third and fourth years after the shock. This observed transmission of a sector-specific labor demand shock to other sectors would generate the observed positive serial dependence in labor demand shocks identified from aggregate MSA employment data.

In summary, the results support the notion that input-output effects are contributing to the observed persistency in labor demand shocks identified from models based on aggregate employment. The results also suggest that the original BK empirical model may perform better in contexts where employment data can be disaggregated into various industries.

7 Discussion

The role that firms play in recoveries from regional recessions has remained a neglected topic of research. Firms may choose to locate new jobs in regions that are economically depressed, attracted by the surplus of labor and low wages. Although on occasion these labor demand responses have been inferred based on empirical models regional adjustment, as we and others have pointed out, these inferences are based on the assumption that the identified shocks of the model are uncorrelated over time. This is an assumption that is not supported by the data. We show that by adopting a VECM we can address the problem caused by persistency because the empirical model separates the structural parameters of interest from the auxiliary parameters that capture serial dependence in the shocks. This permits us to model the effects of an uncorrelated shock, and thereby better measure the labor demand response of interest. We find that in the U.S. the labor demand response is between two and three times larger than the labor supply response—implying that firms are the

Figure 9: Sectoral employment responses to a -1% labor demand shock to manufacturing and mining. Shocks identified using Bartik shift shares. Error bands represent 90% confidence intervals.

main drivers of recoveries and that worker migration only accounts for a relatively small proportion of the recovery. We also find that the recovery process can be highly protracted, taking in excess of twenty years.

Our main finding of a highly-protracted, largely firm-mediated adjustment process has several interesting implications. First, it has significant implications for policy analysis in which regional equilibration is potentially important, such as local development program evaluation (Kline and Moretti, 2011), education policies (Bartik, 2009), local regulation changes (Greenstone, 2002), the effects of foreign immigration on native workers (Borjas, 2006), and the effects of natural disasters on local economies. Because the labor migration response is limited and the labor demand response is protracted, local policy that operates through the labor demand channel can have large and long-lasting effects on local residents.

Second, there has been some controversy regarding the response of migration to regional differences in unemployment rates in the US (see, e.g., Bewley, 1999, p. 371-372), and our findings may help reconcile the opposing sides in this debate. In important early work, Marston (1985) and Blanchard and Katz (1992) conclude that there is a large and fast labor migration response to regional shocks. Although we use the Blanchard and Katz (1992) model as the basis of our analysis, our results are more consistent with the majority of literature which concludes that migration plays a relatively small role in equilibrating regional unemployment rates, and that labor market shocks can have long-lasting effects on regional labor markets (Partridge *et al.*, 2015). In addition, our results are also consistent with individual-level evidence that workers move only in response to very large monetary incentives (Kennan and Walker, 2011).

Third, the implied levels of labor mobility for the US are much closer to those found in other economies, such as Europe (Decressin and Fatás, 1995), Australia (DeBelle and Vickery, 1999) and New Zealand (Grimes, Maré and Morton, 2009). In these non-US studies, the labor demand shocks appear far less persistent, although a complete examination of the data would be necessary before concluding that regional shocks are uncorrelated over time in these economies. Our results could imply that internal labor mobility is more homogenous across different nations than previously thought.

A Appendix

A.1 Data

County level labor market variables are aggregated up to MSAs based on 2009 MSA definitions. The OMB maintains metro/micropolitan statistical area definitions. Historical definitions are available at <http://www.census.gov/population/metro/data/pastmetro.html>. [accessed February 2016]. Because of changes in county boundaries, Boulder and Greeley, CO MSAs (FIPS 14500 and 24540) are combined with Denver, CO MSA (FIPS 19740) to produce a consistent time series. In addition, BLS data omit New Orleans-Metairie-Kenner, LA (FIPS 35380) in some years, and so this MSA is excluded. After these modifications we are left with 363 MSAs from the original 366 MSAs given under the 2009 definitions. Anyone interested in obtaining our data may contact the corresponding author.

A.1.1 BEA Employment Data

BEA employment data are obtained from the personal income and employment summary, denoted with the code “CA04”.

A.1.2 Census Population Data

Up to 2010 population data are obtained from the intercensal population estimates. From 2010 to 2012 we use postcensal population figures. The census constructs intercensal population numbers out of estimates of the components of changing population at the subnational level. They use indicators of births, deaths, international immigration, and internal migration to construct estimates of changes in county populations over time. These data are benchmarked between decennial censuses. To form consistent time series, the only potential age groups are 15-64, total population, 15+, or less than 64 years old. We select 15-64 because this age range most closely represents the working-age population.

A.1.3 LAUS Data

Source data for the LAUS data set include the Current Population Survey (CPS), the Current Employment Statistics survey (CES), the Unemployment Insurance (UI) system, and the Quarterly Census of Employment and Wages (QCEW). The estimation methodology uses as “signal plus noise” approach (see BLS, 2002, and references therein).

The LAUS data set contains a level break in the year 2000. Based on discussions with the BLS, we learned that the incorporation of new inputs and methodological changes associated a 2005 redesign of the LAUS program were carried back only to 2000, creating a break in the series between December 1999 and January 2000. (The redesign is discussed at <http://www.bls.gov/lau/lausredesignqa.htm>.) Because of this, we have adjusted the LAUS data. To do so, we have assumed that the levels of each variable in December 1999 were the same as in January 2000. Growth rates for the preceding months in 1999 are used to obtain levels for each of the other months in that year. These levels are averaged to generate the adjusted 1999 data. From this, the pre-1999 levels of each variable are computed recursively using annual growth rates computed from the original data set.

A.1.4 House Prices and Wages Data

The MSA-level wages we use are also derived from the local area employment and wages data that the BEA produces. Annual wages and salaries are computed by dividing total annual wages (collected at the county level and aggregated to the MSA level) by wage and salary employment. The Federal Housing Finance Administration (FHFA) is the source of data on housing prices. The FHFA house price index (HPI) is not available at the county level; however, the FHFA maintains a consistent time series based on current MSA definitions. Unfortunately, for some MSAs, the FHFA housing data are split into metropolitan divisions (MDs). Because the data represent indices, they cannot be simply aggregated as can labor force, population and salary data. Thus, for MSAs split into MDs, we simply take the house price index from the largest MD. MSA definitions are matched by hand to the 2009 MSA definitions used in this paper, to ensure continuity. There are fewer HPI MSAs than MSAs in the other data sets that we use, as the FHFA requires 1,000 total transactions in the MSA or MD before data are published.

A.1.5 WRLURI Data

The Wharton School of the University of Pennsylvania. These data are provided at a smaller geographic resolution than MSAs, and are aggregated to MSAs using weights provided. See Gyourko *et al.* (2008) for more information about these data.

A.1.6 Bartik Shift-Share Variables

The Bartik (1991) shift-shares are often used as a labor demand shock. Let $E_{i,j,t}$ denote the number of employed persons in location i and industry j in time period t . From these data, we compute two components: National employment growth by industry, and local employment shares by industry. To do so, we follow Saks (2008): To minimize potential endogeneity, when we compute the national employment growth in industry j component of location i 's Bartik variable, we exclude location i 's employment. Thus, we compute national industry employment growth separately for each location: $\Delta e_{j,t}^i = \log(E_{j,t} - E_{i,j,t}) - \log(E_{j,t-1} - E_{i,j,t-1})$ where $E_{j,t} = \sum_{\text{locations } i} E_{i,j,t}$. MSA employment shares are given by $\pi_{i,j,t}^* = E_{i,j,t} / \sum_{\text{industries } j} E_{i,j,t}$. Also following Saks (2008), we subtract national employment growth from local industry employment growth, $\Delta e_t = \log(E_t) - \log(E_{t-1})$ where $E_t = \sum_{\text{industries } j, \text{ locations } i} E_{i,j,t}$. Then the Bartik variable is computed as weighted industry growth, assuming that local employment grows at the same (relative) rate as national industry level employment, $\varepsilon_{i,t}^{d,\text{BT}} = \sum_{\text{industries } j} \pi_{i,j,t} (\Delta e_{j,t}^i - \Delta e_t)$. We use the 3-digit NAICS industry employment data from the Quarterly Census of Employment and Wages (QCEW) from the Bureau of Labor Statistics. We do not use BEA employment since BEA employment by industry and county data are not presented at the same level of disaggregation as the QCEW data.

A.1.7 MSA-by-Industry Data set

Industry data are derived from the BLS QCEW NAICS-based data files, which represent aggregate county industry employment for all establishments that are covered by unemployment insurance. Any suppressed or missing industry-time-county data cause deletion of the entire industry-county observation. MSA data are derived from remaining counties for each industry. Because suppression at the county level is severe for detailed NAICS industries, a relatively broad industry grouping

(nine industries) is used. From these nine industries we generate six sectoral aggregates by adding together Mining and Manufacturing, and Professional, Finance and Other Services (“Professional services”). We aggregate to make to analysis tractable, and two ensure that the various sectors are roughly equal in terms of employment share across time and MSAs. We leave construction as a separate sector due to its volatile and heterogenous behavior.

A.2 Bootstrap

Our bootstrap procedure is as follows. We use 1000 replications.

1. Estimate the reduced-form VECM given in (9) (or (10) in the case of recursive identification) and obtain reduced-form residuals. We employ five lags when estimating the models since the ensuant efficiency losses are outweighed by the potential costs of a more parsimonious but misspecified model (Berkowitz and Kilian, 2000).
2. We re-sample the residuals with replacement, and maintaining cross-sectional and cross-equation correlation in the re-sampled panels by re-sampling only in the time series dimension. We do this in order to preserve any weak-form correlation in the residual that would otherwise cause us to underestimate the variance of the estimator. For example, it is likely that shocks are correlated across locations both according to geography as well as according to industry mix. We reconstruct the data using the reduced-form estimates obtained in step 1. We initialize the simulated data using $\{\Delta Y_{i,1}, \dots, \Delta Y_{i,p}, Z_{i,p}\}_{i=1}^n$, and for models that include the Bartik shift-shares, the same sequence of $\left\{\varepsilon_{i,t}^{d,BT}\right\}_{i,t=p+1}^{n,T}$ are used in each replication since the shift-shares are treated as exogenous.
3. Re-estimate the VECM for each replication sample and compute the impulse response functions of employment, population and labor force growth, and employment rates and participation rates, to a -1% labor demand shock.
4. Reported 90% confidence intervals for the impulse response functions are based on bootstrapped estimates of the standard errors (see Efron, 1981, pp. 139-140). Let $\hat{\sigma}_h$ denote the

bootstrapped standard error of the IRF $\hat{\psi}_h$. Then the confidence interval around $\hat{\psi}_h$ is given by $\left[\hat{\psi}_h - 1.96\hat{\sigma}_h, \hat{\psi}_h + 1.96\hat{\sigma}_h\right]$.

A.3 Model Extensions and Robustness

Figure 10: Responses to a serially uncorrelated -1% labor demand shock using Bartik shift shares. Model specification includes wages and house prices. Error bands represent 90% confidence intervals.

Figure 11: Responses to a serially uncorrelated -1% labor demand shock using Bartik shift-shares. 1st WRLURI Quartile (most elastic WRLURI). Error bands represent 90% confidence intervals.

Figure 12: Responses to a serially uncorrelated -1% labor demand shock using Bartik shift-shares. 4th WRLURI Quartile (most inelastic WRLURI). Error bands represent 90% confidence intervals.

Figure 13: Responses to a serially uncorrelated -1% labor demand shock using Bartik shift shares and using direct estimates of migrants. Error bands represent 90% confidence intervals.

Figure 14: Responses to a serially uncorrelated -1% labor demand shock using recursive identification. Error bands represent 90% confidence intervals.

Figure 15: Responses to a serially uncorrelated -1% labor demand shock using recursive identification. Model specification includes establishments. Error bands represent 90% confidence intervals.

Figure 16: Responses to a serially uncorrelated -1% labor demand shock using recursive identification. Model specification includes wages and house prices. Error bands represent 90% confidence intervals.

Figure 17: Responses to a serially uncorrelated -1% labor demand shock using recursive identification. Employment and working-age population data span 1970-2012. Error bands represent 90% confidence intervals.

Figure 18: Responses to a serially uncorrelated -1% labor demand shock using recursive identification. 1st WRLURI Quartile (most elastic WRLURI). Error bands represent 90% confidence intervals.

Figure 19: Responses to a serially uncorrelated -1% labor demand shock using recursive identification. 4th WRLURI Quartile (most inelastic WRLURI). Error bands represent 90% confidence intervals.

Figure 20: Responses to a serially uncorrelated -1% labor demand shock using recursive identification. Model specification based on direct estimates of migrants. Error bands represent 90% confidence intervals.

Figure 21: Responses to a serially uncorrelated -1% labor demand shock using Bartik shift-shares. Model estimated using Hayakawa (2009) Instrumental Variables. Error bands represent 90% confidence intervals.

Figure 22: Responses to a serially uncorrelated -1% labor demand shock using recursive identification. Model estimated using Hayakawa (2009) Instrumental Variables. Error bands represent 90% confidence intervals.

References

- [1] Arellano, Manuel (2003), Panel Data Econometrics, New York: Oxford University Press.
- [2] Arellano, Manuel (2015), “Modelling Optimal Instrumental Variables For Dynamic Panel Data Models,” Research in Economics, In Press.
- [3] Baltagi, Badi H. (2013), Econometric Analysis of Panel Data, Chichester, UK: John Wiley and Sons.
- [4] Bailey, Natalia, Sean Holly and M. Hashem Pesaran (2016), “A Two-Stage Approach to Spatio-Temporal Analysis with Strong and Weak Cross-Sectional Dependence,” Journal of Applied Econometrics, Vol. 31 (1), pp. 249-280.
- [5] Bartik, Timothy J. (1991), Who Benefits from State and Local Economic Development Policies?, Kalamazoo, MI: W. E. Upjohn Institute for Employment Research.
- [6] Bartik, Timothy J. (1993), “Who Benefits from Local Job Growth: Migrants or the Original Residents?,” Regional Studies Vol. 27 (4), pp. 297-311.
- [7] Bartik, Timothy J. (2009), “What Proportion of Children Stay in the Same Location as Adults, and How Does This Vary Across Location and Groups?,” W. E. Upjohn Institute Working Paper No. 09-145.
- [8] Berkowitz, Jeremy and Lutz Kilian, (2000), “Recent Developments in Bootstrapping Time Series,” Econometric Reviews, Vol. 19 (1), pp. 1-48.
- [9] Bewley, Truman F. (1999), Why Wages Don’t Fall during a Recession?, Cambridge, MA: Harvard University Press.
- [10] Blanchard, Olivier Jean and Lawrence F. Katz (1992), “Regional Evolutions,” Brookings Papers on Economic Activity, Vol. 23 (1), pp. 1-76.
- [11] Blundell, Richard and Stephen Bond (1998), “Initial Conditions and Moment Restrictions in Dynamic Panel Data Models,” Journal of Econometrics, Vol. 87 (1), pp. 115-143.

- [12] Borjas, George J. (1999), "The Economic Analysis from Immigration," in Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics*, Elsevier, ed. 1, vol. 3, no. 3.
- [13] _____ (2006), "Native Internal Migration and the Labor Market Impact of Immigration," *Journal of Human Resources*, Vol. 41 (2), pp. 221-258.
- [14] Borjas, George J., Richard B. Freeman, Lawrence F. Katz, John Dinardo and John M. Abowd (1997), "How Much Do Immigration and Trade Affect Labor Market Outcomes?," *Brookings Papers on Economic Activity*, Vol. 1997 (1), pp. 1-90.
- [15] BLS (2002), "Design and Methodology," *Current Population Survey Technical Paper 63RV*. Downloaded on 1/11/2013 from <http://www.census.gov/prod/2002pubs/tp63rv.pdf>.
- [16] Chen, Yong and Stuart S. Rosenthal (2008), "Local amenities and life-cycle migration: Do people move for jobs or fun?," *Journal of Urban Economics*, Vol. 64 (3), pp. 517-537.
- [17] Debelle, Guy and James I. Vickery (1999), "Labour Market Adjustment: Evidence on Interstate Labour Mobility," *Australian Economic Review*, Vol. 32 (3), pp. 249-63.
- [18] Decressin, Jörg and Antonio Fatás (1995), "Regional Labor Market Dynamics in Europe," *European Economic Review*, Vol. 39 (9), pp. 1627-1655.
- [19] Efron, Bradley (1981), "Nonparametric Standard Errors and Confidence Intervals," *The Canadian Journal of Statistics*, Vol. 9 (2), pp. 139-158
- [20] Greenstone, Michael (2002), "The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures," *Journal of Political Economy*, Vol. 110 (6), pp. 1175-1219.
- [21] Greenwood, Michael J., Gary L. Hunt and John M. McDowell (1986), "Migration and Employment Change: Empirical Evidence on the Spatial and Temporal Dimensions of the Linkage," *Journal of Regional Science*, Vol. 26 (2), pp. 223-234.
- [22] Greenwood, Michael J. and Gary L. Hunt (1989), "Jobs versus Amenities in the Analysis of Metropolitan Migration," *Journal of Urban Economics*, Vol. 25 (1), pp. 1-16.

- [23] Grimes, Arthur, David C. Maré and Melanie Morten (2009), "Adjustment in Local Labour and Housing Markets," *Australasian Journal of Regional Studies*, Vol. 15 (2), pp. 229-348.
- [24] Gyourko, Joseph, Albert Saiz and Anita A. Summers (2008), "A New Measure of the Local Regulatory Environment for Housing Markets," *Urban Studies*, Vol. 45 (3), pp. 693-729.
- [25] Hayakawa, Kazuhiko (2009), "A Simple Efficient Instrumental Variable Estimator for Panel AR (p) Models when both N and T are Large," *Econometric Theory*, Vol. 25 (3), pp. 873-890.
- [26] Hosseinkouchack, Mehdi and Joakim Westerlund (2016), "Modified CADF and CIPS Panel Unit Root Statistics with Standard Chi-squared and Normal Limiting Distributions," *Forthcoming in The Oxford Bulletin of Economics and Statistics*.
- [27] Kaplan, Greg and Sam Schulhofer-Wohl (2012), "Interstate Migration Has Fallen Less Than You Think: Consequences of Hot Deck Imputation in the Current Population Survey," *Demography*, Vol. 49 (3), pp. 1061-1074.
- [28] John Kennan and James R. Walker (2011), "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, Vol. 79 (1), pp. 211-251.
- [29] Kline, Patrick (2008), "Understanding Sectoral Labor Market Dynamics: An Equilibrium Analysis of the Oil and Gas Field Services Industry," *working paper*.
- [30] Kline, Patrick and Enrico Moretti (2011), "Local Economic Development, Agglomeration Economies and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority," *working paper*.
- [31] Kothari, Siddarth, Itay Saporta-Eksten and Edison Yu (2013), "The (Un)importance of Geographical Mobility in the Great Recession," *Review of Economic Dynamics*, Vol. 16 (3), 553-563.
- [32] Lee, Yoonseok and Peter C.B. Phillips (2015), "Model Selection in the Presence of Incidental Parameters," *Journal of Econometrics*, Vol. 188 (2), pp. 474-489.

- [33] Lkhagvasuren, Damba (2012), “Big Locational Differences in Unemployment Despite High Labor Mobility,” *Journal of Monetary Economics*, Vol. 59 (8), pp. 798-814.
- [34] Mark, Nelson C. (2004), *International Macroeconomics and Finance: Theory and Empirical Methods*, Hoboken, NJ: Wiley-Blackwell.
- [35] Molloy, Raven S., Christopher L. Smith and Abigail Wozniak (2011), “Internal Migration in the United States,” *Journal of Economic Perspectives*, Vol. 25 (3), pp. 173-196.
- [36] Nickell, Stephen (1997), “Unemployment and Labor Market Rigidities: Europe versus North America,” *The Journal of Economic Perspectives*, Vol. 11 (3), pp. 55-74.
- [37] Obstfeld, Maurice and Giovanni Peri (1998), “Regional Non-adjustment and Fiscal Policy,” *Economic Policy*, Vol 13 (26), pp. 205-47.
- [38] Partridge, Mark D. and Dan S. Rickman (1999), “Which comes first, jobs or people? An analysis of the recent stylized facts,” *Economic Letters*, Vol. 64 (1), pp. 117-123.
- [39] _____ (2003), “The Waxing and Waning of Regional Economies: The Chicken-Egg Question of Jobs versus People,” *Journal of Urban Economics*, Vol. 53 (1), pp. 76-97.
- [40] _____ (2006), “An SVAR Model of Fluctuations in U.S. Migration Flows and State Labor Market Dynamics,” *Southern Economic Journal*, Vol. 72 (4), pp. 958-980.
- [41] Partridge, Mark D., Dan S. Rickman, Kamar Ali and M. Rose Olfert (2008), “Employment Growth in the American Urban Hierarchy: Long Live Distance,” *The B.E. Journal of Macroeconomics*, Vol. 8 (1), pp. 1-38.
- [42] Partridge, Mark D., Dan S. Rickman, M. Rose Olfert, Kamar Ali (2012), “Dwindling U.S. Internal Migration: Evidence of Spatial Equilibrium or Structural Shifts in Local Labor Markets?,” *Regional Science and Urban Economics*, Elsevier, Vol. 42 (1-2), pp. 375-388.
- [43] Partridge, Mark D., Dan S. Rickman, M. Rose Olfert and Ying Tang (2015), “When Spatial Equilibrium Fails: Is Place-Based Policy Second Best?,” *Regional Studies*. Vol. 49 (8), pp. 1303-1325.

- [44] Pesaran, M. Hashem (2007), "A Simple Panel Unit Root Test in the Presence of Cross Section Dependence," *Journal of Applied Econometrics*, Vol. 22 (2), pp. 265–312.
- [45] Rowthorn, Robert and Andrew J. Glyn (2006), "Convergence and Stability in U.S. Employment Rates," *B.E. Journal of Macroeconomics*, Vol. 6 (1), pp. 1-43.
- [46] Saiz, Albert (2007), "Immigration and Housing Rents in American Cities," *Journal of Urban Economics*, Vol. 61 (2), pp. 345–371.
- [47] Saks, Raven E. (2008), "Job Creation and Housing Construction: Constraints on Metropolitan Area Employment Growth," *Journal of Urban Economics*, Vol. 64 (1), pp. 178-195.
- [48] Saks, Raven E. and Abigail Wozniak (2011), "Labor Reallocation over the Business Cycle: New Evidence from Internal Migration," *Journal of Labor Economics*, Vol. 29 (4), pp. 697-739.
- [49] Sims, Christopher A. (1980), "Macroeconomics and Reality," *Econometrica*, Vol. 48 (1), pp. 1-48.
- [50] Tani, Massimiliano (2003), "Have Europeans Become More Mobile? A note on Regional Evolutions in the EU: 1988-1997," *Economic Letters*, Vol. 80 (1), pp. 23-30.
- [51] Topel, Robert H. (1986), "Local Labor Markets," *Journal of Political Economy*, Vol. 94 (3), pp. S111-143.
- [52] Zabel, Jeffrey (2012), "Migration, Housing Market, and Labor Market Responses to Employment Shocks," *Journal of Urban Economics*, Vol. 72 (2), pp. 267-284.