Economic Quarterly—Volume 102, Number 3—Third Quarter 2016—Pages 169–192

From Stylized to Quantitative Spatial Models of Cities

Sonya Ravindranath Waddell and Pierre-Daniel Sarte

1. INTRODUCTION

Understanding how and why factors of production locate within and around urban areas has been compelling social scientists for at least 150 years. Within mainstream economics, urban economists have been developing modern theories of city systems at least since the 1960s. However, modeling spatial interactions is highly complex, and, therefore, the theoretical literature on economic geography has necessarily focused on stylized settings. For example, a model may have a central business district—where firms are assumed to be located—surrounded by a symmetric circle or on a symmetric line. As the population grows, the scarcity of land prevents consumers (who are also workers) from all settling close to the center, so people move out to where commuting costs are higher but housing costs are lower.

In the models of new economic geography (NEG), urban economists have incorporated advances developed in industrial organization, international trade, and economic growth to remove technical barriers to modeling cities. The field of NEG was initiated primarily by three authors: Fujita (1988), Krugman (1991), and Venables (1996), who all use general equilibrium models with some version of monopolistic competition. The NEG models have been useful in helping to pin down preferences, technology, and endowments and have provided

We thank Daniel Schwam and Daniel Ober-Reynolds for helpful research assistance. The views expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Richmond or the Federal Reserve System. We thank Urvi Neelakantan, Santiago Pinto, and John Weinberg for helpful comments and suggestions.

DOI: http://doi.org/10.21144/eq1020301

some fundamental theoretical explanation for the uneven distribution of economic activity, for multiple equilibria in location choices, and for a small (possibly temporary) asymmetric shock across sites to generate a large permanent imbalance in the distribution of economic activities. However, these models have also imposed structure that is not necessarily evident in the data, and the limitation of the analysis to stylized spatial settings has not enabled an empirical literature that could directly corroborate the theory. In other words, the stylized models have only guided empirical estimation in a way that is divorced from the structure of those models, resulting in empirical research that has been devoid of strong structural interpretations.

More recently, the introduction of quantitative models of international trade (in particular Eaton and Kortum [2002]) have served to develop a framework that connects closely to the observed data. This research does not aim to provide a fundamental explanation for the agglomeration of economic activity but instead aims to provide an empirically relevant quantitative model. This article describes the progression from a simple canonical model of NEG to its counterpart in the quantitative spatial framework. Section 2 engages the literature to develop and understand the progression from the stylized models of the NEG literature to the quantitative spatial models. Section 3 walks through a version of the stylized model, with a linear monocentric city. Section 4 introduces its counterpart as a quantitative spatial model as was laid out in Redding and Rossi-Hansberg (forthcoming). Section 5 provides an example of how the spatial model can be matched to detailed microdata that describe actual interactions in the city. Section 6 concludes.

2. LITERATURE REVIEW

The standard monocentric model of cities came out of a history of work to model spatial allocations. The prototype for understanding how factors of production distribute themselves across land, and how prices govern that distribution, was developed by Johann Heinrich von Thünen in the mid-nineteenth century to describe the pattern of agricultural activities in preindustrial Germany. Von Thünen's model includes an exogenously located marketplace in which all transactions regarding final goods must occur and the differences in land rent and use are determined predominantly by transport costs (Fujita and Thisse 2002). The von Thünen model was both formalized mathematically and enhanced in the second half of the twentieth century—including the formalization of bid-rent curves by William Alonso in his basic urban land model. This basic urban model includes a monocentric city with a center, home to the central business district (CBD), where all jobs are located. The space surrounding the CBD is assumed to be homogenous with only one spatial characteristic: its distance to the CBD. Both the work of von Thünen and that of Alonso depended upon the monocentricity of production activities—i.e., the models rely on one CBD (or market) with surrounding land used for residential (or agricultural) purposes.

Although many early models assumed the existence of the CBD, later work formalized mechanisms for the agglomeration forces that create concentrations of economic activity. The models of NEG, as summarized in Fujita et al. (1999), Fujita and Thisse (2002), and Ottaviano and Thisse (2004), create the framework to explain the imbalance in the distribution of economic activity and better understand how a small shock can generate that imbalance. These NEG models went a long way toward overcoming the fundamental problem that kept economic geography and location theory at the periphery of mainstream economic theory for so long: regional specialization and trade cannot arise in the competitive equilibrium of an economy with homogenous space. This spatial impossibility theorem is discussed more thoroughly in Ottaviano and Thisse (2004) and articulated mathematically in Fujita and Thisse (2002).

Important ideas underlie the development of the NEG models. These ideas (as described in Ottaviano and Thisse [2004]) include that the distribution of economic activity is the outcome of a trade-off between various forms of increasing returns and different mobility costs; price competition, high transport costs, and land use foster the dispersion of production and consumption, and, therefore, firms are likely to cluster in large metropolitan areas when they sell differentiated products and transport costs are low. Cities provide a wide array of final goods and specialized labor markets that make them attractive to consumers/workers, and agglomeration is the outcome of cumulative processes involving both the supply and demand sides. The contribution of NEG was to link those ideas together in a general equilibrium framework with imperfect competition. Some of the earliest work in NEG came from Krugman (1991), who developed a model that showed that the emergence of an industrialized "core" and an agricultural "periphery" pattern depends on transportation costs, economies of scale, and the share of manufacturing in national income (i.e., in consumption expenditures). More specifically, in his model, lower transportation costs, a higher manufacturing share, or stronger economies of scale will result in the concentration of manufacturing in the region that gets a head start compared to other regions. Venables (1996) wrote a model where imperfect competition and transport costs create forward and backward linkages between industries in different locations. He finds that even without labor mobility, agglomeration can be generated through the location decisions of firms in industries that are linked through an input-output structure. The models above develop an argument for agglomeration into a single center of activity. However, other NEG models, most notably Fujita and Ogawa (1982) and Lucas and Rossi-Hansberg (2002), introduced nonmonocentric models where businesses and housing can be located anywhere in the city. The latter models constitute a first step toward building frameworks that more accurately capture the heterogeneity in economic activity across space.

Unfortunately, although the theoretical work on NEG has been relatively rich, the empirical research has been comparatively less rich; establishing causality and controlling for confounding factors has proved challenging in the empirical realm. One challenge, as articulated by Redding and Rossi-Hansberg (forthcoming), is that the complexity of the theoretical models has limited the analysis to stylized spatial settings, such as a few locations, a circle, or a line, and the resulting empirical research has been primarily reduced form in nature. As a result, it is difficult to provide a structural interpretation of the estimated coefficients, and the empirical models cannot either withstand the Lucas critique (coefficients might change with different policy interventions) or necessarily generalize to more realistic environments.

Empirical work, such as the spatial model laid out in Section 4, has been instructed by another field of economics. Developments in the international trade literature have offered mechanisms for better modeling the distribution of economic activity across urban areas. Eaton and Kortum (2002) developed a model of international trade that captures both the comparative advantage that encourages trade and the geographic barriers that inhibit it (e.g., transport costs, tariffs and quotas, challenges negotiating trade deals, etc.). They use the model to solve for the the world trading equilibrium and examine its response to policies.

This framework from the trade literature—combined with the availability of increasingly more granular data—enabled the emergence of new quantitative spatial models in urban economics in which one can carry out general equilibrium counterfactual policy exercises. In addition to offering methodological insights and a mechanism for policy analysis, these quantitative spatial models have made substantive contributions that borrow from, and contribute to, the theoretical literature. For example, Redding and Sturm (2008) provide evidence for a causal relationship between market access and the spatial distribution of economic activity. They show that the division of Germany after

World War II led to a sharp decline in population growth in West German cities close to the new border relative to other West German cities and that this decline was more pronounced for small cities than for large cities. As another example, models such as those developed in Ahlfeldt et al. (2015) and Monte et al. (2016), allow for heterogenous gradients of economic activity within cities that can be matched directly to microdata and that can only be approximated in models such as Fujita and Ogawa (1982) and Lucas and Rossi-Hansberg (2002).

The next section walks through a canonical monocentric urban model and highlights key features that made that model attractive for thinking about the distribution of economic activity across space. In particular, this urban model allows many of a city's features to be endogenous, including its size, population, employment, wages, and commercial land rents. In addition, at different locations within the city, residential population, residential prices, and the consumption of housing services can also be endogenous. In this model, as in the average city, production is concentrated at the center, where the CBD is located, rent gradients decline with distance from the CBD, and population density tends to decrease away from the city center.

3. A STYLIZED MODEL OF CITIES

We consider a linear monocentric city with locations defined on the interval $[-B, B]$, where ℓ denotes the distance from the city center. Each location ℓ is endowed with one unit of land available either for residential housing or production. This analysis focuses on residential localization decisions, i.e., the decisions of households rather than firms.

The Central Business District

All production takes place at the city center, $\ell = 0$, which defines the CBD. Production per unit of land is given by

$$Y = A(L)L^{\beta}, \tag{1}$$

where L denotes labor input and $A(L)$ denotes a production externality. For simplicity, let $A(L) = AL^{\alpha}$, $\alpha < 1 - \beta < 1$, and denote the wage paid to workers by w. This condition ensures that labor demand, L, is decreasing in the wage, w. There exists a unit mass of firms (assuming firms are small and do not internalize the externality) where the representative firm solves

$$\max_{L} A(L)L^{\beta} - wL.$$

It follows that

$$\beta A(L)L^{\beta-1} = w \Leftrightarrow L = \left(\frac{A\beta}{w}\right)^{\frac{1}{1-\alpha-\beta}}. \tag{2}$$

We assume a competitive market with free entry so that in equilibrium firms obtain zero profits. Therefore, the commercial bid rent faced by firms in the business district is

$$q^b = (1-\beta)A^{\frac{1}{1-\alpha-\beta}}\left(\frac{\beta}{w}\right)^{\frac{\alpha+\beta}{1-\alpha-\beta}}. \tag{3}$$

Residential Areas

Workers live in the city at different locations, $\ell \in [-B, B]\backslash\{0\}$, and commute to the city center. Workers who reside at ℓ consume goods, $c(\ell)$, housing services, $h(\ell)$, and experience a commuting cost, $\kappa(\ell) \in [1, \infty)$, that reduces the utility derived from housing and increases with distance from the CBD. In particular, the utility of a worker commuting from location ℓ to the CBD is given by $s\left(\frac{c(\ell)}{\gamma}\right)^{\gamma}\left(\frac{h(\ell)}{(1-\gamma)\kappa(\ell)}\right)^{1-\gamma}$, where $\gamma \in (0,1)$ and s is a service amenity conferred by the city. This approach to modeling commuting costs departs somewhat from the more traditional approach of assuming that disposable income (thus consumption of housing and nonhousing goods) declines with distance from the CBD. In this case, similar to Ahlfeldt et al. (2015), commuting costs enter the utility function multiplicatively, which, as they note, is isomorphic to a formulation in terms of a reduction in effective units of labor. Commuting costs are then ultimately proportional to wages in the indirect utility function.

Conditional on living at location ℓ, a worker then solves

$$u(\ell) = \max_{c(\ell),h(\ell)} s\left(\frac{c(\ell)}{\gamma}\right)^{\gamma}\left(\frac{h(\ell)}{(1-\gamma)\kappa(\ell)}\right)^{1-\gamma},$$
$$\gamma \in (0,1)$$

subject to $c(\ell) + q^r(\ell)h(\ell) = w$,

where $q^r(\ell)$ is the price of a unit of residential housing services at location ℓ. Hence, we have that

$$c(\ell) = \gamma w, \tag{4}$$

$$h(\ell) = \frac{(1-\gamma)w}{q^r(\ell)}, \tag{5}$$

and

$$u(\ell) = s\,[w]^{\gamma} \left[\frac{w}{\kappa(\ell) q^{r}(\ell)} \right]^{1-\gamma}$$
$$= sw\,[\kappa(\ell) q^{r}(\ell)]^{\gamma-1}\,. \tag{6}$$

The Residential Market

Let $\overline{u}$ denote the utility available to workers from residing in alternative cities. To the extent that workers can move to or from another city and are free to reside at any location within the city, it must be the case that in equilibrium $u(\ell) = \overline{u}\ \forall \ell \in [-B, B]$. Therefore, from equation (6), we have that, for any location ℓ,

$$sw\,[\kappa(\ell) q^{r}(\ell)]^{\gamma-1} = sw\,[\kappa(B) q^{r}_{B}]^{\gamma-1}\,, \tag{7}$$

where q^{r}_{B} is the price of land at the boundary of the city defined by the opportunity cost of land at that location, such as an agricultural land rent. Rewriting equation (7) gives residential land rents at different locations within the city,

$$q^{r}(\ell) = \frac{\kappa(B)}{\kappa(\ell)} q^{r}_{B}, \tag{8}$$

where $\frac{\kappa(B)}{\kappa(\ell)} \geq 1\ \forall \ell \in [-B, B]$, since $\kappa(\ell)$ increases with distance from the city center. Thus, residential land rents are highest near the CBD and decrease toward the boundaries of the city as commuting becomes more expensive. However, as seen from equation (5), total housing expenditures in this framework are constant across all locations in the city since $q^{r}(\ell)h(\ell) = (1-\gamma)w$, where $(1-\gamma)$ then represents the income share of housing expenditures.

Recall that each location $\ell \in [-B, B]$ is endowed with one unit of land available for housing. Let $R(\ell)$ denote the residential population living at ℓ. We assume that all available land in the city is fully developed and used by residents. Then, equilibrium in the housing sector requires that

$$R(\ell)h(\ell) = 1. \tag{9}$$

In addition, the residential population living at different locations ℓ in the city must sum up to the supply of labor working in the CBD,

$$\int_{-B}^{B} R(\ell)d\ell = L$$

$$\Rightarrow \int_{-B}^{B} \frac{1}{h(\ell)}d\ell = L \tag{10}$$

Solving for the City Equilibrium

We now describe the city equilibrium, first solving for equilibrium wages as a function of the model parameters, from which all other city allocations immediately follow.

Given equations (5) and (8), equation (10) becomes

$$\int_{-B}^{B} \frac{1}{h(\ell)}d\ell = \int_{-B}^{B} \frac{q^r(\ell)}{(1-\gamma)w}d\ell = \frac{\kappa(B)q_B^r}{(1-\gamma)w}\int_{-B}^{B} \frac{1}{\kappa(\ell)}d\ell = L, \tag{11}$$

which defines the boundaries of the city, $B(L, w)$, as a function of its population and wages given the model's parameters.

Consider for instance the simple symmetric case where $\kappa(\ell) = e^{\kappa|\ell|}$ so that $\kappa(0) = 1$ and $\kappa(B) = e^{\kappa B} > 1$. Then, $\int_{-B}^{B} \frac{1}{\kappa(\ell)}d\ell$ gives

$$\int_{-B}^{B} \frac{1}{\kappa(\ell)}d\ell = \int_{-B}^{B} e^{-\kappa|\ell|}d\ell = 2\int_{0}^{B} e^{-\kappa\ell}d\ell = 2(\frac{-e^{-\kappa\ell}}{\kappa}|_0^B) = \frac{2}{\kappa}(1-e^{-\kappa B}),$$

so that equation (11) becomes

$$\frac{\frac{2}{\kappa}e^{\kappa B}q_B^r(1-e^{-\kappa B})}{(1-\gamma)w} = \frac{\frac{2}{\kappa}q_B^r(e^{\kappa B}-1)}{(1-\gamma)w} = L$$

$$\Rightarrow e^{\kappa B} = 1 + \frac{\kappa(1-\gamma)wL}{2q_B^r}.$$

Using the labor demand equation in equation (2), conditional on the model parameters, the boundaries of the city may then alternatively be expressed in terms of wages only,

$$\begin{aligned} e^{\kappa B} &= 1 + \frac{\kappa(1-\gamma)w}{2q_B^r}\left(\frac{A\beta}{w}\right)^{\frac{1}{1-\alpha-\beta}} \\ &= 1 + \frac{\kappa(1-\gamma)(A\beta)^{\frac{1}{1-\alpha-\beta}} w^{\frac{-\alpha-\beta}{1-\alpha-\beta}}}{2q_B^r}. \end{aligned}$$

Using this last expression, we can solve for equilibrium wages in the city as a function of the model parameters only. Specifically, note that residents' utility at the boundary is given by

$$\overline{u} = sw\left[q_B^r + \frac{\kappa(1-\gamma)(A\beta)^{\frac{1}{1-\alpha-\beta}} w^{\frac{-\alpha-\beta}{1-\alpha-\beta}}}{2}\right]^{\gamma-1}, \tag{12}$$

which defines $w^* = w(s, \kappa, \gamma, A, \alpha, \beta, q_B^r, \overline{u})$.

Proposition 1: *There exists a unique w^* that solves equation (12).*

Proof: *Define* $f(w) = sw\left[q_B^r + \frac{\kappa(1-\gamma)(A\beta)^{\frac{1}{1-\alpha-\beta}} w^{\frac{-\alpha-\beta}{1-\alpha-\beta}}}{2}\right]^{\gamma-1}$. *Then* $\lim_{w\to 0} f(w) = 0$, $\lim_{w\to\infty} f(w) = \infty$, *and, since $f(w)$ is continuous in w, there exists w^* such that $\overline{u} = f(w^*)$. Moreover, since $f(w)$ is strictly increasing in w, w^* is unique.*

Given w^*, all other allocations in the city then immediately follow. In particular, as mentioned in the proposition, given parameter restrictions, the RHS of equation (12) is increasing in w so that w^* then increases with $\overline{u}$. Thus, as the reservation utility from living elsewhere, $\overline{u}$, increases, the city population, $L^* = \left(\frac{A\beta}{w^*}\right)^{\frac{1}{1-\alpha-\beta}}$, falls as residents leave the city, and its boundaries, $B^* = \frac{1}{\kappa}\log(1 + \frac{\kappa(1-\gamma)(A\beta)^{\frac{1}{1-\alpha-\beta}} w^{\frac{-\alpha-\beta}{1-\alpha-\beta}}}{2q_B^r})$, shrink.

The stylized model described above is rich enough to allow for many of a city's features to be endogenous, including its size, population, employment, wages, and commercial land rents. In addition, at different locations within the city, residential population, residential prices, and the consumption of housing services can also be endogenous. These allocations are such that there exists a very direct link between commuting costs to the CBD and residential prices. Specifically, taking equation (8) and using the functional form for commuting costs described above, we can derive a simple expression for the elasticity of residential prices with respect to commuting costs.

Proposition 2: *The elasticity of residential prices with respect to commuting costs,* $\varepsilon_{q^r,\kappa}$, is given by $\kappa(B - |\ell|)$.

Proof:

$$q^r(\ell) = \frac{e^{\kappa B}}{e^{\kappa|\ell|}} q_B^r = e^{\kappa(B-|\ell|)} q_B^r$$

$$\varepsilon_{q^r,\kappa} = \frac{\partial q^r(\ell)}{\partial \kappa} \cdot \frac{\kappa}{q^r(\ell)} = (B - |\ell|) e^{\kappa(B-|\ell|)} q_B^r \cdot \frac{\kappa}{e^{\kappa(B-|\ell|)} q_B^r} = \kappa(B - |\ell|).$$

The proposition above highlights the effect of commuting costs on prices; specifically, this effect is mitigated as we move away from the employment center and is zero at the boundary. Intuitively, away from the city center, residential prices become increasingly pinned down by the agricultural land rent rather than economic activity near the center.

Despite its richness, the stylized model we have just described imposes a number of restrictions on the structure of the city, including its monocentric nature with all production being concentrated in the CBD. Furthermore, residential prices decline monotonically as one moves away from the city center, and there exists a general symmetry and an evenness in allocations and prices across space. This smooth and symmetric aspect of the city is illustrated in Figure 1. In that figure, residential population is highest near the CBD, where the commute is relatively cheap, and decreases monotonically away from the center with the fewest workers living near the boundaries of the city.

In practice, of course, economic activity is more unevenly distributed across space. For example, Figure 2 shows that the city of Richmond, Virginia, has multiple employment clusters, one indeed in the center of the city but two others to the south and west.

This activity reflects a balance of agglomeration forces (e.g., production externalities) and dispersion forces (e.g., commuting costs) that play out in intricate and interrelated ways across space and that lead to substantial variations in allocations and prices across a city. For example, production may take place in different parts of the city so that cities with multiple production centers are not uncommon. In fact, some productive activity potentially takes place at every location in the city. Moreover, residential prices, even if they tend to fall away from a central point in the city, seldom fall monotonically with distance from that center. Instead, residential rents can exhibit substantial variation across locations within the city. This variation reflects the potential complexity of linkages within the city where, for example, the resident population at a given location may depend on the entire distribution of wages offered across the city. Thus, in the next section, we show how

Figure 1 Allocations and Prices in the Monocentric Model

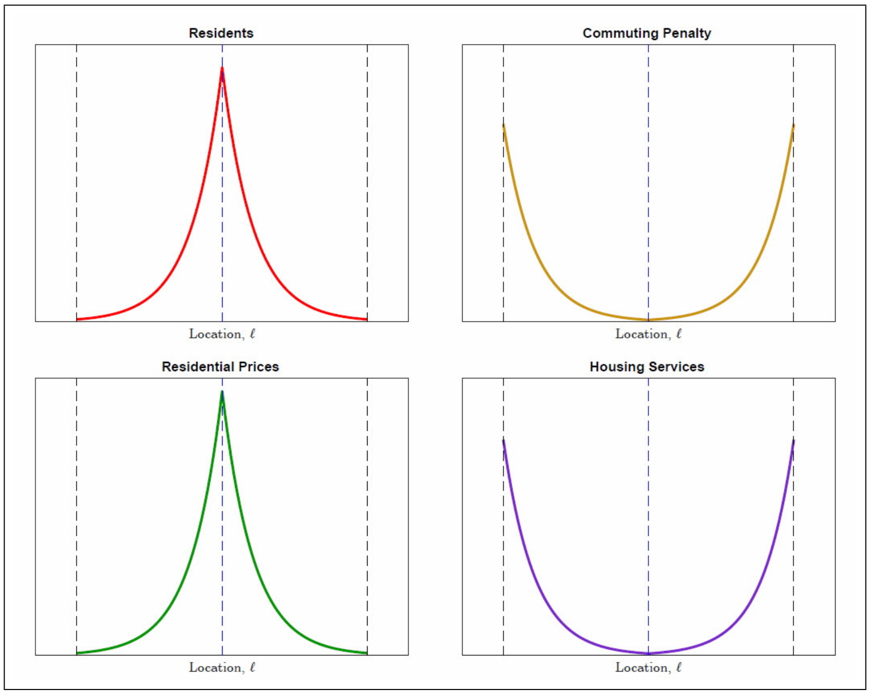

to modify the stylized model presented in this section in a way that can quantitatively account for the spatial allocations and prices it is meant to study.

4. A QUANTITATIVE SPATIAL MODEL OF CITIES

In this section, we show how to adapt the stylized model of the previous section to allow for the heterogeneity in spatial allocations and prices that is typically observed in cities. In doing so, we preserve the basic assumptions on preferences, technology, and endowments of our stylized model to keep the frameworks comparable. Instead of thinking of the city as located on an interval $[-B, B]$, we will think of the city as composed of J distinct locations, indexed by $j \in \{1, ..., J\}$ (or i). In the mapping to data, these locations may represent city blocks, census tracts, or counties. It is this key change that will allow us to ensure that the model is at least able to match given observed spatial allocations of, for example, resident population, land rents, employment, or wages across locations in a city. Any subsequent counterfactual exer-

Figure 2 Density of Primary Jobs

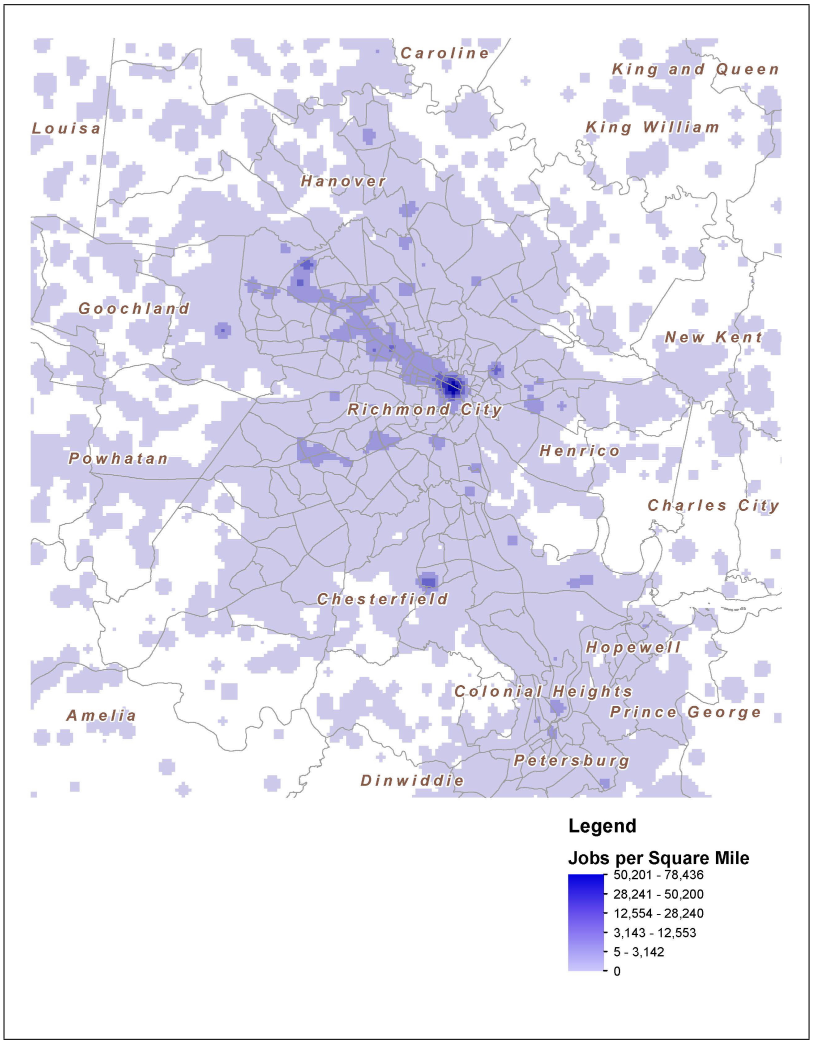

cise involving a change to some exogenous aspect of the city is then grounded in a model that is able to exactly replicate uneven spatial observations that reflect, at least in part, complex linkages between decisions involving where to reside and where to work within the city. For example, the model would enable us to understand the effect of a new urban policy, such as one that provides housing assistance or subsidized transportation.

In a model where every location could potentially be used for both residential and production purposes, a central component of a quantitative spatial model is the representation and matching of distinct pairwise commuting flows from any location j in the city to any other location i. This step will rely on an approach developed by Eaton and Kortum (2002) in modeling trade flows between locations. As in the model developed in the previous section, this analysis will focus on residential localization decisions, i.e., the decisions of households rather than firms. Unlike in the previous model, the commercial bid rent schedule is nondegenerate and reflects variations in productivity and wages across locations.

Firms

Production per unit of land in the business district of each location i is given by

$$Y_i = A(L_i)L_i^{\beta}, \tag{13}$$

analogously to equation (1), where L_i denotes labor input and $A(L_i)$ denotes a production externality that we assume is local (so only employment in i affects the productivity of businesses in i). For simplicity, let $A(L_i) = A_i L_i^{\alpha}$, $\alpha < 1 - \beta < 1$, and denote the wage paid to workers in location i by w_i. There exists a unit mass of firms (assuming that firms are small and do not internalize the externality) where the representative firm solves

$$\max_{L_i} A(L_i)L_i^{\beta} - w_i L_i.$$

It follows that

$$\beta A(L_i)L_i^{\beta-1} = w \Leftrightarrow L_i = \left(\frac{A_i \beta}{w_i}\right)^{\frac{1}{1-\alpha-\beta}}. \tag{14}$$

As in the previous model, firms operate in a competitive market with free entry and thus obtain zero profits in equilibrium. The implied commercial bid rent schedule faced by firms in the business district is

$$q_i^b = (1-\beta)A_i^{\frac{1}{1-\alpha-\beta}} \left(\frac{\beta}{w_i}\right)^{\frac{\alpha+\beta}{1-\alpha-\beta}}. \tag{15}$$

Note the similarities between equations (14) and (15), and the analogous equations in the previous section, equations (2) and (3).

Residents

In each location j of the city, there exists a residential area composed of a continuum of residents who commute to the business areas of different locations i for work. These residents differ in their preferences for where to work in the city according to a random idiosyncratic component s. Unlike the previous model where s was a city amenity distributed uniformly across locations, in this model, s is an individual-specific preference component. Conditional on living in a particular location j, this preference component captures the idea that residents of j may have idiosyncratic reasons for commuting to different locations i in the city. We model the idiosyncratic preference component associated with residing in location j and working in location i as scaling the utility of the residents of region j, where s is drawn from a Fréchet distribution specific to that particular commute,

$$F_{ij}(s) = e^{-\lambda_{ij} s^{-\theta}}, \ \lambda_{ij} > 0, \ \theta > 0. \tag{16}$$

Residents of j who commute to i incur an associated cost, $\kappa_{ij} \in [1, \infty)$, that, analogous to the previous section, reduces the utility derived from housing. Thus, conditional on living in j and working in i, the problem of a resident having drawn idiosyncratic utility s is given by

$$u_{ij}(s) = \max_{c_{ij}(s), h_{ij}(s)} \ s \left(\frac{c_{ij}(s)}{\gamma_j} \right)^{\gamma_j} \left(\frac{h_{ij}(s)}{(1-\gamma_j)\kappa_{ij}} \right)^{1-\gamma_j},$$

$$\gamma_j \in (0,1)$$

subject to $c_{ij}(s) + q_j^r h_{ij}(s) = w_i$,

where q_j^r is the price of a unit of residential housing services at location j. Hence, we have that

$$c_{ij}(s) = \gamma_j w_i, \tag{17}$$

$$h_{ij}(s) = \frac{(1-\gamma_j) w_i}{q_j^r}, \tag{18}$$

and

$$\begin{aligned} u_{ij}(s) &= s \left[w_i \right]^{\gamma_j} \left[\frac{w_i}{\kappa_{ij} q_j^r} \right]^{1-\gamma_j} \\ &= s w_i \left[\kappa_{ij} q_j^r \right]^{\gamma_j - 1}. \end{aligned} \tag{19}$$

Note the similarities between equations (17), (18), and (19) and the analogous equations in the previous sections, equations (4), (5), and (6).

Aggregation

The setup we have just described allows for a considerable degree of heterogeneity within the city compared to the stylized model presented earlier. In particular, all locations allow for simultaneous use by both businesses and residents (mixed use), individuals living in any location may commute to any other location for work, and commute costs between any two locations are specific to that pair of locations, so that it is possible to take into account the particular geographical makeup or road infrastructure of a city. However, having allowed for this high level of heterogeneity in the city, it becomes important to be able to aggregate economic activity at the level of a location, such as a census tract for practical purposes. The steps in this subsection address this question.

Distribution of Utility

Since residents of j who work in i have different preferences s, drawn from equation (16), for commuting to that location, it follows that

$$G_{ij}(u) = \Pr(u_{ij} < u) = F_{ij}\left(\frac{u\left[\kappa_{ij} q_j^r\right]^{1-\gamma_j}}{w_i}\right),$$

or

$$G_{ij}(u) = e^{-\Phi_{ij} u^{-\theta}},\ \Phi_{ij} = \lambda_{ij} w_i^{\theta}\left[\kappa_{ij} q_j^r\right]^{(\gamma_j - 1)\theta}. \tag{20}$$

Each resident of j chooses to commute to the location i that offers maximum utility of all possible locations. Therefore,

$$\begin{aligned} G_j(u) &= \Pr(\max_i\{u_{ij}\} < u) = \prod_i \Pr(u_{ij} < u) \\ &= \prod_i e^{-\Phi_{ij} u^{-\theta}}. \end{aligned}$$

Thus, it follows that

$$G_j(u) = e^{-\Phi_j u^{-\theta}},\ \Phi_j = \sum_i \Phi_{ij}. \tag{21}$$

In other words, the distribution of resident utility in each location j of the city is itself a Fréchet distribution. The expected utility from residing in j is then given by

$$u_j = \Gamma\left(\frac{\theta-1}{\theta}\right)\left(q_j^r\right)^{\gamma_j - 1}\left(\sum_i \lambda_{ij} w_i^{\theta} \kappa_{ij}^{(\gamma_j-1)\theta}\right)^{\frac{1}{\theta}}, \tag{22}$$

where $\Gamma(.)$ is the Gamma function.[1] The expected utility from living in j, therefore, is a weighted average of the utility gained from commuting to the different business areas (raised to the θ). Observe that in contrast to utility in the stylized model, equation (6), the expected utility from living in location j of the city now involves not only the price of housing at that location, but also information about the entire city, including the entire distribution of wages and associated commuting costs, since residents of j can in principle commute to any other location i to work.

Commuting Patterns

Let π_{ij} represent the proportion of residents living at location j and commuting to location i. Commuting patterns can then be described by the following relationship,

$$R_{ij} = \pi_{ij} R_j,$$

where R_{ij} and R_j are, respectively, the number of residents commuting from j to i and the total number of residents living at j. In particular,

$$\pi_{ij} = \Pr\left[u_{ij} > \max_{n \neq i}\{u_{nj}\}\right].$$

From equation (20), we have that $G_{ij}(u) = e^{-\Phi_{ij}u^{-\theta}}$ so that $g_{ij}(u) = \theta u^{-(\theta+1)}\Phi_{ij}e^{-\Phi_{ij}u^{-\theta}}$. It follows that

$$\pi_{ij} = \int_0^\infty \theta u^{-(\theta+1)}\Phi_{ij}e^{-\Phi_{ij}u^{-\theta}}\widetilde{G}_j(u)du, \tag{23}$$

where $\widetilde{G}_j(u)$ is defined as in equation (21) but with $\widetilde{\Phi}_j = \sum_{n \neq j} \Phi_{nj}$, which also implies that $\Phi_j = \widetilde{\Phi}_j + \Phi_{ij}$. In Appendix B, we show that this expression reduces to

$$\pi_{ij} = \frac{\lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j - 1)\theta}}{\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j - 1)\theta}}. \tag{24}$$

In other words, the proportion of residents living in j and commuting to i for work depends on wages earned in i adjusted for commuting costs when coming from j, relative to a weighted average of wages earned elsewhere adjusted for the corresponding commute (raised to the θ).

[1] A derivation of this result is given in Appendix A.

The residential price at location j does not affect commuting patterns from j to i since it is specific to j and is faced by any resident who wants to live at j regardless of commute. By construction, $\sum_i \pi_{ij} = 1$.

The Residential Market

Recall that $h_{ij}(s) = h_{ij}$ represents housing consumption for those living in j and commuting to i. It follows that average housing per resident at location j, h_j, is given by

$$\begin{aligned} h_j &= \sum_i \pi_{ij} h_{ij} \\ &= \frac{(1-\gamma_j)}{q_j^r} \sum_i \pi_{ij} w_i. \end{aligned}$$

As in the stylized model of the previous section, we assume that each location is endowed with one unit of land available for housing and that this land is fully developed.[2] In equilibrium, therefore, the residential market must satisfy $R_j h_j = 1$ similarly to equation (9) or

$$R_j = \frac{q_j^r}{(1-\gamma_j) \sum_i \pi_{ij} w_i}. \tag{25}$$

As in the previous section, let $\overline{u}$ denote the utility available to individuals from residing in alternative cities. To the extent that workers can move to or from another city, and are free to reside at any location within the city, it must be the case that in equilibrium $u_j = \overline{u}\ \forall j$. Therefore, we have that for any location,

$$\begin{aligned} \overline{u} &= \Gamma\left(\frac{\theta-1}{\theta}\right) \left(q_j^r\right)^{\gamma_j - 1} \left(\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}\right)^{\frac{1}{\theta}} \\ \Rightarrow q_j^r &= \left[\frac{\overline{u}}{\Gamma\left(\frac{\theta-1}{\theta}\right)}\right]^{\frac{1}{\gamma_j-1}} \left(\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}\right)^{\frac{1}{\theta(1-\gamma_j)}}. \end{aligned} \tag{26}$$

Comparing the residential price at location j, q_j^r, with its simpler analog in the stylized model in equation (8), it is clear that the quantitative spatial model allows residential prices to be determined by many more

[2] Owens et al. (2017) present a more flexible model in which residential land in any one location may be vacant, partially developed with some areas left for developers to build on, or fully developed. In that model, a coordination problem arises between developers and residents (no one wants to be the first mover) that potentially traps neighborhoods in an equilibrium where they remain vacant.

factors, including the distribution of wages across all locations in the city as well as all commuting costs. It is this richness that allows for spatial variation in allocations and prices across locations in the city that is unavailable in the more stylized framework of the previous section.

The City Labor Market

Since $\pi_{ij}R_j$ denotes the number of residents living in location j who commute to the business area of location i for work, labor market clearing in the city requires that

$$L_i = \sum_{j=1}^{J} \pi_{ij} R_j,$$

or alternatively

$$\left(\frac{A_i \beta}{w_i}\right)^{\frac{1}{1-\alpha-\beta}} = \sum_{j=1}^{J} \pi_{ij} R_j. \tag{27}$$

Solving for the City Equilibrium

We represent the parameters of the quantitative spatial model in a vector, $\mathcal{P} = (\alpha, \beta, \theta, \overline{u},\ \gamma_j, A_i, \kappa_{ij}, \lambda_{ij})$. Conditional on $\mathcal{P}$, equations (24), (25), (26), and (27) then make up a system of $J^2 + 3J$ equations in the same number of unknowns, $\pi_{ij}(\mathcal{P})$, $R_j(\mathcal{P})$, $q_j^r(\mathcal{P})$, and $w_i(\mathcal{P})$.

Importantly, the equilibrium allocations in this model allow for considerably more heterogeneity than in the stylized model of the previous section. Since they are specific to locations within the city, equilibrium allocations of the quantitative spatial model such as commuting patterns, $\pi_{ij}(\mathcal{P})$, or equilibrium prices, such as residential prices, $q_j^r(\mathcal{P})$, and wages, $w_i(\mathcal{P})$, may be directly matched to their data counterpart at the block or census tract level. In contrast, equilibrium allocations of the stylized model in the previous section could only be indexed by distance, ℓ, from a central point in the city. The next section addresses this last point in more detail.

Unlike the conventional monocentric model of the previous section, equilibrium existence and uniqueness are more challenging to prove in a quantitative spatial framework. However, Appendix C summarizes the key equations needed to compute the model equilibrium and provides an algorithm that yields the corresponding numerical solution given the model's parameters, $\mathcal{P}$. Moreover, despite its added complexity, the quantitative spatial model retains some degree of analytical tractability.

For instance, as in the monocentric model of the previous section, we can derive a simple expression for the elasticity of residential prices with respect to commuting costs.

Proposition 3: *The elasticity of residential prices with respect to commuting costs*, $\varepsilon_{q_j^r,\kappa_{ij}}$, is given by $-\pi_{ij}$.

Proof: We have that $q_j^r = \left(\frac{\bar{u}}{\Gamma\left(\frac{\theta-1}{\theta}\right)}\right)^{\frac{1}{\gamma_j-1}} \left(\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}\right)^{\frac{1}{\theta(1-\gamma_j)}}$. Then

$$\frac{\partial q_j^r}{\partial \kappa_{ij}} =$$

$$\left(\frac{\bar{u}}{\Gamma\left(\frac{\theta-1}{\theta}\right)}\right)^{\frac{1}{\gamma_j-1}} \frac{1}{\theta(1-\gamma_j)} \left(\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}\right)^{\frac{1}{\theta(1-\gamma_j)}-1} (\gamma_j-1)\theta\lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta-1}$$

$$= -1 \cdot \frac{\lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta-1}}{\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}} q_j^r$$

It follows that

$$\varepsilon_{q_j^r,\kappa_{ij}} = \frac{\partial q_j^r}{\partial \kappa_{ij}} \cdot \frac{\kappa_{ij}}{q_j^r} = -\frac{\lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta-1} q_j^r}{\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}} \cdot \frac{\kappa_{ij}}{q_j^r} = -\pi_{ij}.$$

This finding is intuitive. A 1 percent increase in commute costs between any two locations, κ_{ij}, lowers residential prices by the share of residents affected by that commute. In this relatively simple spatial environment, even if it allows for more flexibility than the monocentric setup, the relationship is exact. More importantly, unlike the analogous elasticity in the more stylized model, the share of residents is itself an endogenous outcome that depends on all of the city's characteristics, $\mathcal{P}$, and thus will move along with the entire distribution of wages and population across locations in any policy experiment.

5. MATCHING THE QUANTITATIVE SPATIAL MODEL TO URBAN MICRODATA

As elaborated upon in earlier sections, it is now possible to model cities by matching these types of quantitative spatial models to available microdata. For the purpose of the discussion below, the parameters in $\mathcal{P}$ fall into two broad classifications: citywide parameters and location- or neighborhood-specific parameters. The parameters in $\mathcal{P}$ that are not location specific have generally accepted values in the literature. For example, Monte et al. (2016) estimate θ (the parameter that governs the shape of the distribution of the idiosyncratic preference, s, of commuting from i to j) to be 4.43. Ciccone and Hall (1996) estimate α (the production externality) to be 0.06, and Ahlfeldt et al. (2015) estimate β (the parameter that defines the relationship between labor and output, separate from the externality) to be 0.80, while $\overline{u}$ is a normalizing constant. The parameters that are location-specific potentially present a greater computational challenge since there are many of them. For example, in a city with $1,000$ census tracts, there are $1,000,000$ λ_{ij}'s. Other location-specific parameters, such as pairwise commuting costs, κ_{ij}, may be directly calibrated to data on distances or commuting times.

The Longitudinal Employer-Household Dynamics Origin-Destination Employment Statistics provide reliable data on cities at the census tract level, including commuting patterns (π_{ij}), resident population (R_j), employment (L_i), and wages (w_i). Other detailed data on cities are also available; for example, residential prices (q_j^r) are available from CoreLogic or county assessors' offices. In general, such data show considerable unevenness across space within a city.

We now describe how, in our simple framework, the location-specific parameters of our quantitative spatial model may be calibrated to exactly match, in equilibrium, all pairwise commuting patterns, π_{ij}, the exact distribution of population across space, R_j, and thus also the distribution of employment, L_i, and the exact distribution of wages in the city, w_i, in a given benchmark period. In particular, we choose the parameters of the model $(\lambda_{ij}, \gamma_j, A_i)$ to match the observations (π_{ij}, R_j, w_i) as equilibrium outcomes. In this way, counterfactual exercises involving a change to some exogenous aspect of the city, or a change in urban policy, are rooted in a model that, as a benchmark, is able to exactly match basic observed allocations and prices in the city as equilibrium outcomes.

Recall that commuting patterns, π_{ij}, are given by equation (24),

$$\pi_{ij} = \frac{\lambda_{ij} w_i^{\theta} \kappa_{ij}^{(\gamma_j - 1)\theta}}{\sum_i \lambda_{ij} w_i^{\theta} \kappa_{ij}^{(\gamma_j - 1)\theta}},$$

where $\kappa_{ij} \in [1, \infty)$. If $\pi_{ij} = 0$, then either $\lambda_{ij} = 0$ or $\kappa_{ij} \rightarrow \infty$. Commuting patterns can be alternatively expressed in terms of the Head and Ries (2001) index,

$$\frac{\pi_{ij}}{\pi_{jj}} = \frac{\lambda_{ij} w_i^{\theta} \kappa_{ij}^{(\gamma_j - 1)\theta}}{\lambda_{jj} w_j^{\theta} \kappa_{jj}^{(\gamma_j - 1)\theta}}.$$

Then, conditional on $\theta = 4.43$ and values for γ_j, the preference parameters, λ_{ij}, can then be chosen to be consistent with commuting patterns,

$$\lambda_{ij} = \pi_{ij} \left(\frac{w_j}{w_i}\right)^{\theta} \left(\frac{\kappa_{ij}}{\kappa_{\min}}\right)^{(1-\gamma_j)\theta} \left(\frac{\lambda_{jj}}{\pi_{jj}}\right), \tag{28}$$

where $\kappa_{\min}$ is a lower bound on κ_{jj}.[3] Since κ_{ij} may be directly inferred from commuting costs data, this approach to obtaining λ_{ij} to match commuting patterns presumes we are also able to match wages, w_i, as part of the model inversion. We show below that this can indeed be done through the choice of location-specific productivities, A_i. With this in mind, we first choose γ_j so as to match the distribution of resident population across space, R_j, conditional on π_{ij} and w_i.

The number of residents in location j is given by

$$R_j = \frac{q_j^r}{(1-\gamma_j)\sum_i \pi_{ij} w_i}. \tag{29}$$

Using equations (26) and (28) with $\kappa_{\min} = 1$, and the normalization $\left(\frac{\lambda_{jj}}{\pi_{jj}}\right) = 1$, for all locations j, equation (29) simplifies to

$$R_j(\gamma_j) = \left(\frac{\Gamma\left(\frac{\theta-1}{\theta}\right) w_j}{\bar{u}}\right)^{\frac{1}{1-\gamma_j}} \frac{1}{(1-\gamma_j)\sum_i \pi_{ij} w_i}. \tag{30}$$

Notice that as $\gamma_j \to 0^+$, $R_j(\gamma_j) \to \frac{\Gamma\left(\frac{\theta-1}{\theta}\right) w_j}{\bar{u}\sum_i \pi_{ij} w_i}$ and as $\gamma_j \to 1^-$, $R_j(\gamma_j) \to \infty$. Therefore, one may choose $\bar{u}$ so that $R_j > \frac{\Gamma\left(\frac{\theta-1}{\theta}\right) w_j}{\bar{u}\sum_i \pi_{ij} w_i}$ for all j and numerically solve the expression in equation (30) to obtain a set of γ_j that exactly matches the distribution of R_j, conditional on π_{ij} and w_i. Since the distribution of γ_j then depends on $\bar{u}$, and γ_j represents the share of income spent on housing in a given census tract j, we choose $\bar{u}$ so that the mean of γ_j is 0.76 as in Ahlfeldt et al. (2015).

Since commuting patterns π_{ij} can be exactly matched, given γ_j and wages w_i, through the choice of λ_{ij} in equation (28), it remains only to ensure that the model is consistent with the spatial distribution of wages in an equilibrium benchmark version of the model. Using equation (26), we can write the city labor market clearing condition, equation (27) as

$$\left(\frac{A_i \beta}{w_i}\right)^{\frac{1}{1-\alpha-\beta}} = \sum_{j=1}^{J} \pi_{ij} R_j,$$

[3] Since $\sum_i \pi_{ij} = 1$, one needs to also normalize the λ_{ij}'s, for example, $\frac{\lambda_{jj}}{\kappa_{jj}} = 1 \ \forall j$.

in which case we can simply choose location-specific productivities, A_i, to ensure that equilibrium benchmark wages exactly match the distribution of wages in the city,

$$A_i = \frac{w_i}{\beta} \left\{ \sum_{j=1}^{J} \pi_{ij} R_j \right\}^{1-\alpha-\beta}. \tag{31}$$

Observe that on the righthand side of equation (31), we are free to use the data on commuting patterns, π_{ij}, and residential population, R_j, since those are matched by construction through the choices of λ_{ij} and γ_j in equations (28) and (30).

6. CONCLUSION

The development of the new quantitative equilibrium models has initiated a more robust and realistic framework with which to model cities. This framework will enable urban economists to provide empirically driven insight into future theoretical or structural work on how cities grow, shrink, and change. By offering a more accurate grounding for empirical models, it will also allow for more robust counterfactual policy exercises that can inform practitioners and policymakers regarding strategies for urban development.

REFERENCES

Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf. 2015. "The Economics of Density: Evidence from the Berlin Wall." *Econometrica* 83 (November): 2127–89.

Ciccone, Antonio, and Robert E. Hall. 1996. "Productivity and the Density of Economic Activity." *American Economic Review* 86 (March): 54–70.

Eaton, Jonathan, and Samuel Kortum. 2002. "Technology, Geography, and Trade." *Econometrica* 70 (September): 1741–79.

Fujita, Masahisa. 1988. "A Monopolistic Competition Model of Spatial Agglomeration: Differentiated Product Approach." *Regional Science and Urban Economics* 18 (February): 87–124.

Fujita, Masahisa, Paul Krugman, and Anthony J. Venables. 1999. *The Spatial Economy: Cities, Regions, and International Trade.* Cambridge, Mass.: MIT Press.

Fujita, Masahisa, and Hideaki Ogawa. 1982. "Multiple Equilibria and Structural Transition of Non-Monocentric Urban Configurations." *Regional Science and Urban Economics* 12 (May): 161–96.

Fujita, Masahisa, and Jacques-François Thisse. 2002. *Economics of Agglomeration: Cities, Industrial Location, and Regional Growth.* Cambridge: Cambridge University Press.

Krugman, Paul. 1991. "Increasing Returns and Economic Geography." *Journal of Political Economy* 99 (June): 483–99.

Lucas, Robert E., and Esteban Rossi-Hansberg. 2002. "On the Internal Structure of Cities." *Econometrica* 70 (July): 1445–76.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg. 2016. "Commuting, Migration, and Local Employment Elasticities." Princeton University Working Paper (October).

Ottaviano, Gianmarco, and Jacques-François Thisse. 2004. "Agglomeration and Economic Geography." In *Handbook of Regional and Urban Economics*, vol. 4, ed. J. Vernon Henderson and Jacques-François Thisse. North Holland: Elsevier: 2563–608.

Owens, Raymond III, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte. 2017. "Rethinking Detroit." Working Paper 23146. Cambridge, Mass.: National Bureau of Economic Research. (February).

Redding, Stephen J., and Esteban Rossi-Hansberg. Forthcoming. "Quantitative Spatial Economics." *Annual Review of Economics*.

Redding, Stephen J., and Daniel M. Sturm. 2008. "The Costs of Remoteness: Evidence from German Division and Reunification." *American Economic Review* 98 (December): 1766–97.

Venables, Anthony J. 1996. "Equilibrium Locations of Vertically Linked Industries." *International Economic Review* 37 (May): 341–59.

APPENDIX: APPENDIX A

Under the maintained assumptions, expected utility at location j is given by

$$u_j = \int_0^\infty \theta \Phi_j u^{-\theta} e^{-\Phi_j u^{-\theta}} du.$$

Consider the change in variables,

$$y = \Phi_j u^{-\theta},\ dy = -\theta \Phi_j u^{-(\theta+1)} du.$$

Then, we have that

$$u_j = \int_0^\infty \Phi_j^{\frac{1}{\theta}} y^{\frac{-1}{\theta}} e^{-y} dy = \Gamma\left(\frac{\theta-1}{\theta}\right) \Phi_j^{\frac{1}{\theta}},$$

from which equation (22) in the text follows.

APPENDIX: APPENDIX B

From equation (23), we have that

$$\begin{aligned} \pi_{ij} &= \int_0^\infty \theta u^{-(\theta+1)} \Phi_{ij} e^{-\Phi_{ij} u^{-\theta}} e^{-\widetilde{\Phi}_j u^{-\theta}} du \\ &= \int_0^\infty \theta u^{-(\theta+1)} \Phi_{ij} e^{-\Phi_j u^{-\theta}} du. \end{aligned}$$

Consider the change of variables,

$$y = \Phi_j u^{-\theta},\ dy = -\theta \Phi_j u^{-(\theta+1)} du.$$

It then follows that

$$\pi_{ij} = \Phi_{ij} \int_0^\infty \theta u^{-(\theta+1)} e^{-\Phi_j u^{-\theta}} du$$
$$= \frac{\Phi_{ij}}{\Phi_j} \int_0^\infty e^{-y} dy$$
$$= \frac{\Phi_{ij}}{\Phi_j},$$

where recall that $\Phi_{ij} = \lambda_{ij} w_i^\theta \left[\kappa_{ij} q_j^r\right]^{(\gamma_j - 1)\theta}$ and $\Phi_j = \left(q_j^r\right)^{(\gamma_j-1)\theta} \sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}$. Equation (24) in the text directly follows.

APPENDIX: APPENDIX C

The Basic Set of Equations and Unknowns

Let L_i^D and L_i^S represent, respectively, labor demand and labor supply in location i. Given a benchmark or counterfactual set of parameters, $\mathcal{P}$, each endogenous variable in the model can ultimately be expressed as depending only on a vector of wages across all locations, $\boldsymbol{w} = (w_1, ..., w_J)'$, and $\mathcal{P}$,

$$\pi_{ij}(\boldsymbol{w}) = \frac{\lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}}{\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}},$$
$$q_j^r(\boldsymbol{w}) = \left[\frac{\bar{u}}{\Gamma\left(\frac{\theta-1}{\theta}\right)}\right]^{\frac{1}{\gamma_j-1}} \left(\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}\right)^{\frac{1}{\theta(1-\gamma_j)}},$$
$$R_j(\boldsymbol{w}, \boldsymbol{\pi}, \boldsymbol{q}) = \frac{q_j^r}{(1-\gamma_j)\sum_i \pi_{ij} w_i},$$
$$L_i^D(\boldsymbol{w}) = \left(\frac{A_i \beta}{w_i}\right)^{\frac{1}{1-\alpha-\beta}},$$
$$L_i^S(\boldsymbol{\pi}, \boldsymbol{R}) = \sum_j \pi_{ij} R_j.$$

Then, finding an equilibrium of the model is equivalent to finding a vector of wages that clears the labor market; $L_i^D(\boldsymbol{w}) = L_i^S(\boldsymbol{w})$ for all locations $i = 1, ..., J$. Put another way, the task is to find a vector $\boldsymbol{w}^* \in \mathbb{R}_+^J$ such that

$$(L_i^D - L_i^S)(\boldsymbol{w}^*) =$$

$$\left(\frac{A_i\beta}{w_i}\right)^{\frac{1}{1-\alpha-\beta}} - \sum_j \frac{\lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}}{\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}} \frac{\left[\frac{\bar{u}}{\Gamma\left(\frac{\theta-1}{\theta}\right)}\right]^{\frac{1}{\gamma_j-1}} \left(\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}\right)^{\frac{1}{\theta(1-\gamma_j)}}}{(1-\gamma_j)\sum_i \frac{\lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}}{\sum_i \lambda_{ij} w_i^\theta \kappa_{ij}^{(\gamma_j-1)\theta}} w_i}$$

$$= \mathbf{0}$$

Several algorithms exist to numerically solve nonlinear system of equations, and MATLAB's fsolve function handles this particular system well.

Numerical Algorithm

Some quantitative spatial models can result in systems whose features (such as nondifferentiability in the presence of thresholds or binding constraints on available land) make traditional algorithms difficult to apply. In such cases, a simple "guess-and-iterate" method can be constructed to calculate solutions. We outline such a method here as it applies to our model.

1. Choose a tolerance level $\varepsilon > 0$ and guess a vector of wages, $\boldsymbol{w}_n$.
2. Calculate the implied matrix of flows: $\pi_{ij}(\boldsymbol{w}_n) = \frac{\lambda_{ij} w_{n,i}^\theta \kappa_{ij}^{(\gamma_j-1)\theta}}{\sum_i \lambda_{ij} w_{n,i}^\theta \kappa_{ij}^{(\gamma_j-1)\theta}}$.
3. Calculate the implied prices: $q_j^r(\boldsymbol{w}_n) = \left[\frac{\bar{u}}{\Gamma\left(\frac{\theta-1}{\theta}\right)}\right]^{\frac{1}{\gamma_j-1}} \left(\sum_i \lambda_{ij} w_{n,i}^\theta \kappa_{ij}^{(\gamma_j-1)\theta}\right)^{\frac{1}{\theta(1-\gamma_j)}}$.
4. Using the prices and flows calculated in steps two and three, calculate the implied number of residents: $R_j(\boldsymbol{w}_n) = \frac{q_j^r(\boldsymbol{w}_n)}{(1-\gamma_j)\sum_i \pi_{ij}(\boldsymbol{w}_n) w_i}$.
5. Using the residents calculated in step four, calculate the implied labor supply in each labor market: $L_i^S(\boldsymbol{w}_n) = \sum_j \pi_{ij}(\boldsymbol{w}_n) R_j(\boldsymbol{w}_n)$.
6. Calculate the implied labor demand in each labor market: $L_i^D(\boldsymbol{w}_n) = \left(\frac{A_i\beta}{w_i}\right)^{\frac{1}{1-\alpha-\beta}}$.

7. At the vector of wages, $\boldsymbol{w}_n$, calculate the implied excess demand for labor in each market: $L_i^D(\boldsymbol{w}_n) - L_i^S(\boldsymbol{w}_n)$.

8. If the aggregate labor market fails to clear, $\sum_i |L_i^D(\boldsymbol{w}_n) - L_i^S(\boldsymbol{w}_n)| > \varepsilon$, then update the vector of wages as follows:

$$\boldsymbol{w}_{n+1} = \boldsymbol{w}_n + \delta \left(L_i^D(\boldsymbol{w}_n) - L_i^S(\boldsymbol{w}_n) \right),$$

   for some $\delta > 0$. This updating rule raises wages in markets where there is excess demand for labor or reduces it where there is excess supply.

9. Repeat steps two through eight until $\sum_i |L_i^D(\boldsymbol{w}_n) - L_i^S(\boldsymbol{w}_n)| \leq \varepsilon$.

Economic Quarterly—Volume 102, Number 3—Third Quarter 2016—Pages 197–223

Beveridge Curve Shifts and Time-Varying Parameter VARs

Thomas A. Lubik, Christian Matthes, and Andrew Owens

At first glance, many macroeconomic time series exhibit some form of nonlinearity. For instance, output growth and inflation show less volatility in the 1980s and 1990s than during the Great Inflation period of the 1970s, an observation that has been labeled the Great Moderation. Over the business cycle, the unemployment rate exhibits an asymmetric sawtooth pattern whereby it rises rapidly during downturns and declines only gradually during a recovery. Many price variables, such as exchange rates or commodity prices, appear stable for a long period followed by sudden level shifts. The literature has studied various specific forms of nonlinearity—such as structural breaks, time-varying volatility, or business cycle asymmetries—using sophisticated time-series methods ranging from threshold and Markov switching to vector-autoregressions (VARs) with time-varying parameters and stochastic volatility. The result as to whether there is nonlinearity in the data has been mixed.[1] A key issue in this

John Bailey Jones, Daniel Ober-Reynolds, Zhu Wang, and John Weinberg provided helpful comments that improved the quality and exposition of the article. We also wish to thank conference participants at the 2015 Meeting of the Society for Computational Economics in Taipei and the 2015 Symposium of the Society for Nonlinear Dynamics and Econometrics in Oslo for useful feedback and suggestions. The views expressed in this paper are those of the authors and should not necessarily be interpreted as those of the Federal Reserve Bank of Richmond or the Federal Reserve System. Lubik: Research Department, Federal Reserve Bank of Richmond. P.O. Box 27622, Richmond, VA 23261. Email: thomas.lubik@rich.frb.org. Matthes: Research Department, Federal Reserve Bank of Richmond. P.O. Box 27622, Richmond, VA 23261. Email: christian.matthes@rich.frb.org. Owens: Department of Economics, University of Rochester. 280 Hutchison Road, Box 270156, Rochester, NY 14627. Email: owensaa@gmail.com.

[1] See, for instance, Hamilton (1989), Primiceri (2005), Sims and Zha (2006), or Amir-Ahmadi, Matthes, and Wang (2016), who use a wide range of empirical methodolgies, data sets, and sample periods.

DOI: http://doi.org/10.21144/eq1020302

literature is that tests for nonlinearity tend to have low power against linear alternatives.

Against this background, time-varying parameter vector-autoregressions (TVP-VARs) with stochastic volatility have emerged as a promising framework to analyze a wide range of underlying nonlinearities.[2] In this class of models, the coefficients of the time-series representation for economic data are allowed to vary over time. The idea is that this feature approximates the underlying nonlinearity in the data-generating process to a satisfactory degree and in a parsimonious manner. For instance, a structural break in a deep parameter, or a switch in regimes, could be captured by a shock to the innovation in a random-walk VAR coefficient. Since TVP-VARs offer this flexibility, that is, since they can be understood as approximations to a wide range of underlying nonlinear behavior, they have become increasingly popular in recent years as empirical modeling devices.[3] TVP-VARs are estimated almost exclusively using Bayesian methods. This is necessitated by the fact that, as with any model that features many parameters, the use of prior information is crucial to deliver sensible estimates. In TVP-VARs the choice of priors is of special importance because, with standard sample sizes, they have a substantial impact on how much of the variation in observables is attributed to stochastic volatility versus time variation in other coefficients.[4] At the same time, there is a growing sense, e.g., Lubik and Matthes (2015), that the conclusions drawn from the TVP-VAR literature warrant skepticism. More specifically, TVP-VARs often find not much time variation in the lag coefficients. Instead, they attribute the variation seen in the data to movements in volatilities as the right incidence of shocks can in principle capture a range of time-series patterns.

The purpose of this article is to investigate the extent to which an inherently nonlinear TVP-VAR with stochastic volatility does, in fact, pick up nonlinear features in the underlying data. We do so by applying the TVP-VAR methodology to data generated from a simple (but nonlinear) search and matching model that is designed to generate endogenous shifts in parameters. We thus ask whether a TVP-VAR is capable of detecting the resulting nonlinearity in Beveridge curve dynamics. We follow standard procedure and prescriptions in the literature to specify the TVP-VAR and to choose the prior. The results from these benchmark exercises show that the concerns about proper

[2] See Lubik and Matthes (2015) for an introduction and survey of TVP-VARs.

[3] See Canova, Ferroni, and Matthes (2015) for a discussion of these issues.

[4] This point is demonstrated by means of a simple example in Lubik and Matthes (2015).

Figure 1 The Beveridge Curve over the Great Recession

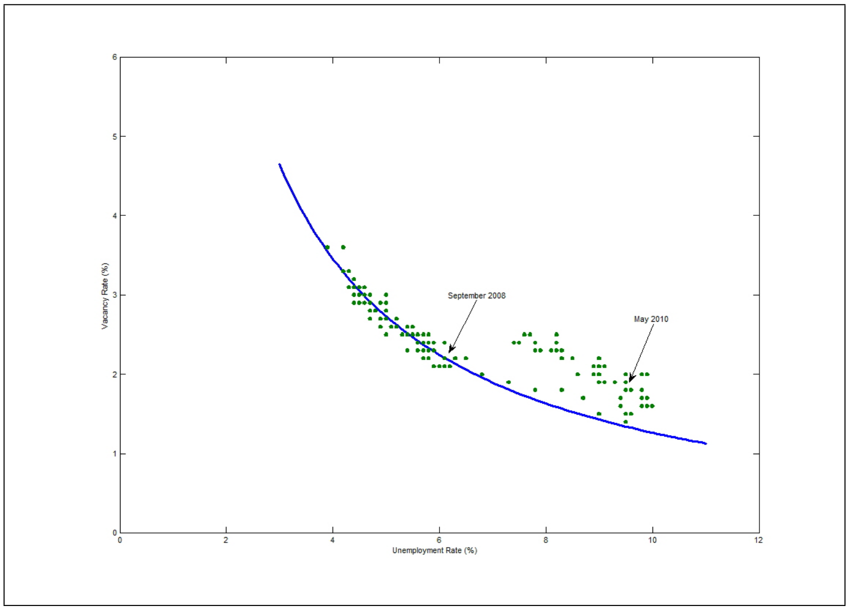

attributions of the sources of nonlinearity are warranted. We attempt to resolve some of these concerns by means of an alternative strategy in choosing priors with only partial success. While these findings are largely negative and are also highly conditional on the chosen theoretical model environment, we argue that they serve as a cautionary tale when conducting and interpreting TVP-VAR studies.

Our chosen framework to analyze these issues is the labor market regularity captured by the so-called Beveridge curve. It describes the joint behavior of unemployment and vacancies over the business cycle and is often seen as indicative of the state of the labor market. The Beveridge curve depicts a negative relationship between these two variables, whereby movements along this curve reflect expansions and recessions. The behavior of the curve over the course of the Great Recession and its aftermath has attracted much interest in the literature (e.g., Barlevy 2011; Lubik 2013; or Şahin et al. 2014). Figure 1 shows the Beveridge curve for data over the Great Recession period. The unemployment-vacancies relationship is often represented by a scatter plot of the two series against each other, resulting in a downward-sloping curve. For purposes of illustration, in Figure 1 we fitted a regression line to data from 2001 up to September 2008, which cluster tightly around the Beveridge curve. At the onset of the Great Recession, the unemployment rate rises rapidly and vacancies fall. In the graph, the data points start moving off the normal curve and appear to settle at a location above their normal, or expected, level. In other words, during the Great Recession, the Beveridge curve appears to have shifted outward in a discrete manner, which could be indicative of a structural break in a labor market parameter.

More generally, the Beveridge curve over the last sixty years reveals a substantial degree of nonlinearity (see Benati and Lubik 2014). There are discernible inward and outward shifts, tilts, and even the occasional slope reversal over short periods. This is not necessarily prima facie evidence of the presence of nonlinearities since these patterns can be rationalized through the right incidence of various shocks (e.g., Blanchard and Diamond 1989; Barlevy 2011; or Lubik 2013). It is nevertheless suggestive of underlying structural changes in the labor market. We take this observation as a starting point for our investigation into the practice of TVP-VAR estimation.

We develop a simple search and matching model of the labor market, where we allow for endogenous threshold switching in a key parameter, namely in the efficiency of a match between employer and job seeker. This match efficiency is captured by a level parameter in the matching function and summarizes the efficacy of the labor market. We assume that it can take on two values, which indicate different but parallel locations of the Beveridge curve. A high level of match efficiency translates into a location of the Beveridge curve closer to the origin, whereby a lower level shifts it outward. Under high efficiency, employers need to open fewer vacancies for the same number of job seekers to fill a desired number of positions. The economy switches between the two efficiency parameters when a threshold embedded in the model is crossed endogenously. We assume that this threshold is given by a low level of output that we associate with a weak labor market performance. This threshold can be reached with a sequence of bad and persistent productivity draws; that is, in this case, the labor market exhibits damage to the extent that the Beveridge curve shifts only when a recession is deep and drawn out. In terms of the behavior of the model, this threshold switch implies nonlinearity in the dynamics of the economic variables.

In order to study the implications of this specific form of nonlinearity for empirical modeling, we solve the full nonlinear model and simulate data on unemployment and vacancies. We then estimate a Bayesian TVP-VAR with stochastic volatility on these data and assess how well the nonlinear *atheoretical* time-series model captures the underlying nonlinearity in the model. Given a standard initialization and choice of priors, the evidence suggests that the TVP-VAR attributes almost all of the changes in the simulated data to changes in the reduced-form innovation variances. We argue that this raises doubts as to the validity of TVP-VAR models with standard priors in detecting shifts. In order to address this shortcoming, we suggest an approach that tries to elicit priors for the TVP-VAR, but it is only moderately successful. In order to better capture the time variation in parameters, researchers will need to adapt the priors to the question at hand in more sophisticated ways. One possibility that delivers better performance is to estimate the hyperparameters associated with the parameters governing the amount of time variation in the model (Amir-Ahmadi, Matthes, and Wang 2017).

The contribution of this article is twofold. First, using simulated data, we study to what extent a generic TVP-VAR with stochastic volatility deals with a specific form of nonlinearity in these underlying data. Our results suggest that some findings of this literature should be regarded with skepticism since they attribute too much of this nonlinearity to time variation in the shocks rather than to structural breaks in the underlying model parameters. Second, and of independent interest, we demonstrate how Beveridge curve shifts can be explained conceptually via an endogenous mechanism that moves the economy between a high-performing and a low-performing labor market. This mechanism can thus be used to address issues like hysteresis, where temporary shocks, such as business cycle shocks, can have permanent effects.

The article is structured as follows. In the next section, we lay out our simple modeling framework of the standard search and matching model and describe how we introduce the threshold-switching mechanism that leads to nonlinearity in the model. The second section describes how we calibrate the model. In this section, we also describe simulation results from the model and discuss the TVP-VAR that we use to estimate the simulated data. Section 3 presents the estimation results and details the shortcomings of the TVP-VAR approach in this environment, while Section 4 introduces an alternative method to elicit priors for the TVP-VAR. The final section concludes.

1. A STRUCTURAL MODEL OF BEVERIDGE CURVE SHIFTS

We now describe the simple structural labor market framework that we use to model the Beveridge curve. We hereby draw heavily from the existing literature, most prominently Shimer (2005). The specification of the model follows Lubik (2013). Our working assumption is that the Beveridge curve has experienced a structural shift, as seen by the evolution of unemployment and vacancies in Figure 1. We model the structural break in terms of a threshold-switching process: when a target variable, aggregate output in our case, hits a threshold in terms of deviations from its long-run level, it triggers a shift in a structural labor market parameter. The idea is to capture the observation that Beveridge curve shifts appear to occur during strong and deep recessions and expansions (see Benati and Lubik 2014).

We assume that in our model economy time is discrete and the time period is a quarter. The labor market in this economy is characterized by search and matching frictions, which help rationalize the existence of equilibrium unemployment. Specifically, a job, that is, a relationship between a worker and a firm for the purpose of engaging in production, is the outcome of a matching process. New jobs M are generated by combining unemployed job seekers U with job openings (vacancies) V. This process can be represented by a constant returns matching function, $M_t = m(s_t)U_t^{\xi}V_t^{1-\xi}$, where $0 < \xi < 1$ is the match elasticity. $m(s_t) > 0$ is the match efficiency that captures the ease with which the unemployed are transformed into workers.

We assume that match efficiency is subject to structural shifts. Specifically, the level parameter in the matching function $m(s_t)$ can switch between two values, $s_t \in \{s_H, s_L\}$, with $m(s_L) < m(s_H)$. In our framework, the switch is generated endogenously by a trigger mechanism, in contrast to the exogenous regime changes in Markov-switching models. We implement this trigger by tying it to the severity of the business cycles. Whenever GDP deviates too much from its current target level, the labor market experiences a structural shift in terms of a change in the matching efficiency. As Lubik (2013) argues, Beveridge curve shifts are most parsimoniously and plausibly modeled by a change in this one parameter. More specifically, one can show that declines in match efficiency are associated with outward shifts of the curve.

For the purposes of capturing Beveridge curve dynamics, we assume that the threshold mechanism is attached to aggregate output. More specifically, we assume that match efficiency $m_t = m(s_t)$ follows a threshold process:

$$m_t = \begin{cases} m(s_H) \text{ if } Y_t \geq \underline{Y} \\ m(s_L) \text{ if } Y_t < \underline{Y} \end{cases}, \text{where } m(s_H) > m(s_L). \tag{1}$$

Y_t is aggregate output and $\underline{Y}$ is the threshold at which the labor market experiences a structural shift. In the simple search and matching framework, we assume linear production so that Y_t is given by:

$$Y_t = A_t N_t, \tag{2}$$

where N_t is the stock of employed workers, and A_t is an aggregate productivity process that obeys the law of motion:

$$\log A_t = (1 - \rho_A) \log \overline{A} + \rho_A \log A_{t-1} + \varepsilon_{A,t}, \tag{3}$$

where $0 < \rho_A < 1$ and $\varepsilon_{A,t} \sim \mathcal{N}(0, \sigma_A^2)$. We normalize the mean of the process $\overline{A}$ to a value of unity without loss of generality.

The dynamics of the model are such that sequences of low and persistent productivity draws—in other words, a recession—will occasionally move aggregate output below the threshold $\underline{Y}$. This damages the labor market in the sense that match efficiency declines and the Beveridge curve shifts outward. This shift is persistent because of the persistence in the productivity process and the inherent persistence of employment in the search and matching framework. Once the recession abates, the labor market recovers in terms of a switch back to a "normal" level of match efficiency. In that sense, our framework shares similarities with the "plucking" model of recessions, where the economy is plucked away occasionally from its normal evolution due to a deep recession but then transitions back over time.

The dynamics of employment are governed by the following relationship:

$$N_t = (1 - \chi_t) \left[N_{t-1} + m(s_{t-1}) U_{t-1}^{\xi} V_{t-1}^{1-\xi} \right]. \tag{4}$$

This is a stock-flow identity that relates the stock of employed workers N to the flow of new hires, $M = mU^{\xi} V^{1-\xi}$, into employment. The timing assumption is such that variations in match efficiency do not affect employment contemporaneously. Unemployment is defined as:

$$U_t = 1 - N_t, \tag{5}$$

where the labor force is normalized to 1. Inflows to unemployment arise from exogenous job destruction at rate $0 < \chi < 1$. We assume that the separation rate χ follows the process:

$$\log \chi_t = (1 - \rho_\chi) \log \overline{\chi} + \rho_\chi \log \chi_{t-1} + \varepsilon_{\chi,t}, \tag{6}$$

where $0 < \rho_\chi < 1$ and $\varepsilon_{\chi,t} \sim \mathcal{N}(0, \sigma_\chi^2)$.

The matching function can be used to define the job-matching rate, i.e., the probability that a firm is matched with a worker:

$$q(\theta_t) = \frac{M_t}{V_t} = m_t \theta_t^{-\xi}, \tag{7}$$

where $\theta_t = V_t / U_t$ is labor market tightness. From the perspective of an individual firm, the aggregate match probability $q(\theta_t)$ is exogenous, and hence new hires are linear in number of vacancies posted for individual firms: $M_{it} = q(\theta_t) V_{it}$.

A firm chooses the optimal number of vacancies V_t to be posted and its employment level N_t by maximizing the intertemporal profit function:[5]

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[A_t N_t - w_t N_t - \kappa V_t \right], \tag{8}$$

subject to the employment accumulation equation (4). Profits are discounted at rate $0 < \beta < 1$. Wages paid to the workers are w, while $\kappa > 0$ is a firm's fixed cost of opening a vacancy. The first order conditions are:

$$N_t \quad : \qquad \mu_t = A_t - w_t + \beta E_t (1 - \chi_{t+1}) \mu_{t+1}, \tag{9}$$
$$V_t \quad : \qquad \kappa = q(\theta_t) \beta E_t (1 - \chi_{t+1}) \mu_{t+1}, \tag{10}$$

where μ_t is the multiplier on the employment equation. Combining these two first-order conditions results in the *job-creation* condition:

$$\frac{\kappa}{q(\theta_t)} = \beta E_t \left[(1 - \chi_{t+1}) \left(A_{t+1} - w_{t+1} + \frac{\kappa}{q(\theta_{t+1})} \right) \right], \tag{11}$$

which captures the trade-off faced by the firm. The marginal, effective cost of posting a vacancy, $\frac{\kappa}{q(\theta_t)}$, that is, the per-vacancy cost κ adjusted for the probability that the position is filled, is weighed against the discounted benefit from the match. The latter consists of the surplus generated by the production process net of wage payments to the workers plus the benefit of not having to post a vacancy again in the next period.

Wages are determined based on the Nash bargaining solution: surpluses accruing to the matched parties are split according to a rule that maximizes the weighted average of the respective surpluses. We relegate the full discussion of the derivation to the Appendix (see also, Lubik 2013). The resulting wage equation is:

$$w_t = \eta \left(A_t + \kappa \theta_t \right) + (1 - \eta) b. \tag{12}$$

Wage payments are a weighted average of the worker's marginal product A_t, of which the worker can appropriate a fraction η, and the outside option b, of which the firm obtains the portion $(1 - \eta)$. Moreover, the presence of fixed vacancy posting costs leads to a hold-up problem where the worker extracts an additional $\eta \kappa \theta_t$ from the firm.

[5] For ease of exposition and notation, we will drop the firm-specific subscripts and discuss the problem of a representative optimizing firm with the understanding that firms are ex-ante heterogeneous in this framework.

We can substitute the wage equation and the job-matching rate into the job-creation condition to obtain:

$$\frac{\kappa}{m_t}\theta_t^{\xi} = \beta E_t\left[\left(1-\chi_{t+1}\right)\left\{\left(1-\eta\right)\left(A_{t+1}-b\right) - \eta\kappa\theta_{t+1} + \frac{\kappa}{m_{t+1}}\theta_{t+1}^{\xi}\right\}\right]. \tag{13}$$

Firms are more willing to post vacancies if productivity shocks increase the wedge to the outside option of the worker; they are less willing if there are expected separations as this will reduce the present value of a hired worker.

In our simulation and empirical analysis, we make use of the simple structure of the model. The dynamics can be fully described by two equations, the employment accumulation equation (4) and the job-creation condition (13), after convenient substitutions. Intuition for why an outward shift of the Beveridge curve is generated by a fall in match efficiency can be gleaned from equation (4) and the logic of the matching function. At any given unemployment rate, firms would need to post more vacancies to achieve a target hiring quota since the matching process is now less efficient.[6] However, there is also a countervailing effect, namely through the influence of match efficiency on firms' vacancy posting decisions. A fall in m raises effective vacancy posting costs as captured by the left-hand side of the job-creation condition (13). This implies that vacancies are increasing in match efficiency. The overall effect of a change in m therefore depends on the interaction of these two margins. As Lubik (2013) shows, the stock-flow identity of the law of motion has to hold in equilibrium, so that the first effect via the matching function dominates and shifts the Beveridge curve outward for a smaller m, whereas the effect via the job-creation margin generates movements along this equilibrium relationship. We now turn to a discussion of our solution and simulation approach.

2. SIMULATION AND ESTIMATION

Calibration

We calibrate our model to representative parameter values in the literature. Our benchmark calibration rests on the parameter estimates

[6] Formally, this can also be seen from the steady-state representation of the employment equation (4), which describes an equilibrium locus of combinations of U and V such that inflows and outflows to (un)employment are balanced:

$$V = \left(\frac{1}{m}\frac{\chi}{1-\chi}\right)^{\frac{1}{1-\xi}}\left(\frac{1-U}{U}\right)^{\frac{1}{1-\xi}}U.$$

Table 1 Calibration

Parameter	Value	Source
Separation Rate $\overline{\chi}$	0.036	Shimer (2005); Monthly JOLTS Data
Match Elasticity ξ	0.49	Beveridge Curve Estimation: Lubik (2013)
Match Efficiency m_H	0.90	Beveridge Curve Estimation: Lubik (2013)
Match Efficiency m_L	0.70	Beveridge Curve Estimation: Lubik (2013)
Benefit b	0.90	Hagedorn and Manovskii (2008)
Bargaining η	0.49	Hosios-Condition: $\eta = \xi$
Job Creation Cost κ	0.18	Imputed from Steady-State Sample Means $\overline{V} = 2.6\%$ and $\overline{U} = 5.2\%$.
Discount Factor β	0.99	Annual Real Interest Rate
Productivity $\overline{A}$	1.00	Normalized
Threshold Value $\underline{Y}$	0.91	Cumulative Decline in U.S. GDP 2008–10
AR(1) Coefficient ρ_A	0.95	Standard Value
AR(1) Coefficient ρ_s	0.95	Standard Value
StD Productivity σ_A	0.01	Standard Value
StD Separation Rate σ_s	0.01	Standard Value

in Lubik (2013) for the period 2000–08, after which a potential shift in the Beveridge curve appears evident from the data (see Figure 1). The calibrated parameters are reported in Table 1. We set the mean of the separation rate to a value of 0.036. This follows the value reported in Shimer (2005) for monthly data. We choose the match efficiency in the high state $m_H = 0.90$ and in the low state $m_H = 0.70$ based on the estimate in Lubik (2013). The match elasticity is set to $\xi = 0.49$. These values broadly determine the slope and the location of the Beveridge curve in a scatter plot of vacancies and unemployment. We set the discount factor $\beta = 0.99$ and choose the bargaining parameter by imposing the Hosios-condition for social efficiency, $\eta = \xi = 0.49$. As mentioned before, we normalize the mean of the level of productivity to $\overline{A} = 1$. Next, we assume that the outside option of the worker makes up 90 percent of the productivity level, $b/\overline{A} = 0.9$. The calibration is therefore close to that of Hagedorn and Manovskii (2008), who argue that a high outside option for the worker is needed to match the cyclical properties of the data. The job-creation condition can then be used to back out the cost parameter κ for a given level of unemployment and vacancies. We compute these from the sample averages for the period 2000–08, $\overline{V} = 2.6$ percent and $\overline{U} = 5.2$ percent. This implies $\kappa = 0.18$. Finally, we set the threshold value for $\underline{Y} = 0.91$ to approximate the cumulative decline in U.S. GDP over the course of the Great

Figure 2 Policy Functions

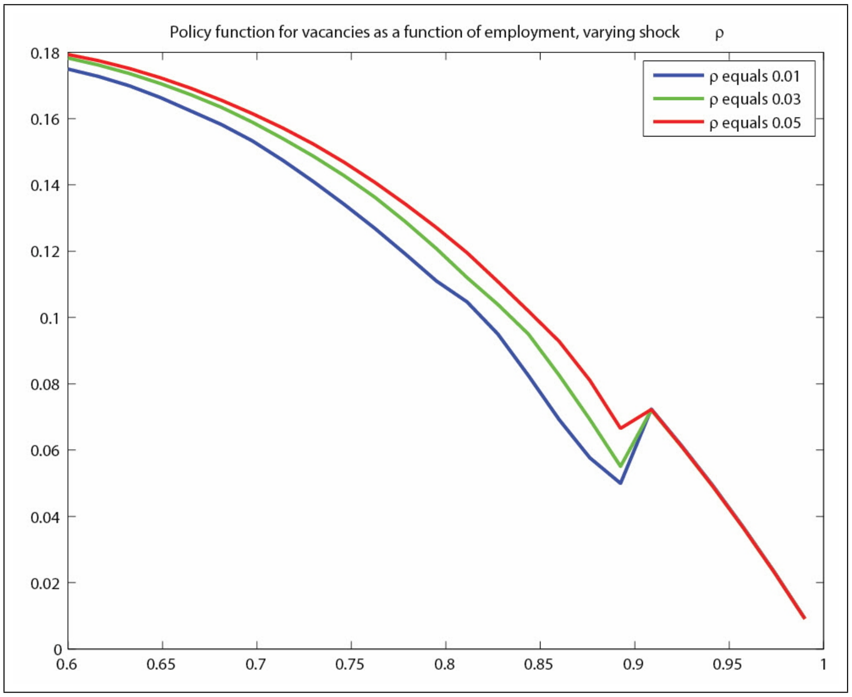

Recession of 2007–09. We set the persistence parameter of the technology process and the separation rate to 0.95 and the standard deviations of the respective innovations to 0.01.

Model Simulation and Discussion

We solve the search and matching model with threshold switching in match efficiency fully nonlinearly by means of the monotone mapping algorithm. The algorithm computes an approximation of the firm's decision rule, which determines the number of vacancy postings given the economy's state variables: employment N_t, the exogenous productivity shock A_t, and the separation rate process χ_t. The algorithm is detailed in the Appendix.

In order to understand the underlying dynamics of the model before we turn to the estimation exercise, we compute the policy functions under the baseline calibration for given realizations of the shocks. The key driving force behind the shifts in the Beveridge curve are movements in productivity. An adverse enough realization of productivity A_t can drive output below the threshold value, which then generates a switch to a lower match efficiency. However, equilibrium outcomes across the threshold and within the distinct regions depend on the subtle interplay between the state variables. To give a sense of the nonlinearities present in our model, we plot the policy function for vacancies in Figure 2. For the purposes of this exercise, we hold the productivity shock fixed at its unconditional mean $\overline{A} = 1$. Vacancies are graphed against the model's sole endogenous state variable, namely the level of employment. We plot this relationship for different realizations of the separation rate $\overline{\chi}$.

Figure 2 shows the key aspect of the model. The policy function has two distinct regions that coincide with the two distinct states of the labor market. For given productivity, the policy function is discontinuous at the implied threshold level, $\overline{N} = \overline{Y} = 0.91$. To the right of the threshold, the labor market is in its normal state with match efficiency $m(s_H)$, and to the left, it has suffered from a deterioration of the latter. We also note that vacancies are decreasing in employment. When employment is high (and unemployment low), few vacancies are being posted because the vacancy-unemployment ratio θ is high and the labor market is tight. That is, the firm's probability of finding a worker is low relative to the costs of hiring him. When employment is low, the labor market is awash with job seekers, so firms can more easily recoup the implicit hiring costs. We also note that the vacancy policy function is increasing in the productivity shock.

The policy function for the high match efficiency case tends to lie to the right and above the respective function in the low efficiency regime. Other things being equal, a lower match efficiency reduces the firm's hiring probability and thereby the incentive to post vacancies relative to the high efficiency scenario. An interesting pattern emerges when we additionally vary the policy function across separation rates. We find that the higher the separation rate, the higher the vacancy postings for given productivity and employment. More separations mean higher churn, so for given employment, more vacancies need to be posted. The differences between the policy functions, however, are quantitatively small for the low efficiency case and almost nonexistent under high efficiency.[7] What is interesting is that the relationship between the separation rate and levels of match efficiency appears nonlinear in its effect on vacancy postings.

[7] This is consistent with the empirical finding in Lubik (2009) and the assumption and interpretation in Shimer (2005) that movements in the separation rate are not key drivers of labor market fluctuations.

Table 2 Selected Moments

	$\sigma(V)$	$\sigma(U)$	$\sigma(V/U)$	$\rho(V,U)$
Sample 1	0.53	1.88	0.32	-0.35
Sample 2	0.49	2.10	0.39	-0.51
Sample 3	0.43	1.22	0.41	-0.47
Sample 4	0.39	0.57	0.41	-0.43
Sample 5	0.35	1.09	0.66	-0.69

We now simulate the model for 590 periods under the benchmark calibration. We discard the first 450 periods as burn-in. We are thus left with a sample of size 140, of which we will use the first forty observations as a training sample in the estimation of the VAR. This leaves us with 100 periods, or twenty-five years, of data for the actual estimation. Table 2 reports moments for five representative samples. We present this as a first pass for whether our regime-switching framework can potentially capture salient labor market facts. The last column shows the correlation between unemployment and vacancies, which is considerably negative and ranges from -0.35 to -0.69 but is below the correlation found in U.S. data. Nevertheless, the model can replicate to some extent the strongly negative comovement between these two labor market variables.

The model is less successful in terms of volatilities. The first two columns of Table 2 report the standard deviations of vacancies and unemployment relative to the standard deviation of (labor) productivity $A_t = Y_t/N_t$ as in Shimer (2005). Vacancies are roughly half as volatile as productivity, while unemployment is twice as volatile for samples 1 and 2. The standard deviation drops considerably in sample 4, while samples 3 and 5 show the volatilities of the driving process and the endogenous variables as roughly equal. The low volatility of vacancies is also reflected in that of labor market tightness V/U. Our framework thus falls prey to the critique espoused in Shimer (2005), namely that the basic search and matching model has difficulty replicating the observed high volatility of unemployment and vacancies. As Lubik (2009) shows, this can be remedied by additional shocks to the model such as the exogenous variations in the separation rate, but this comes at the price of reducing the correlation between U and V since a shock to separations moves unemployment and vacancies in the same direction.

Figure 3 shows data plots of the five simulation samples, including the training sample, in the same order as presented in Table 2. Each row in the graph represents one simulation. The panels on the left show time series plots of unemployment (in red) and vacancies (in blue). The middle column shows the same data as a scatter plot in order to

Figure 3 Simulated Data

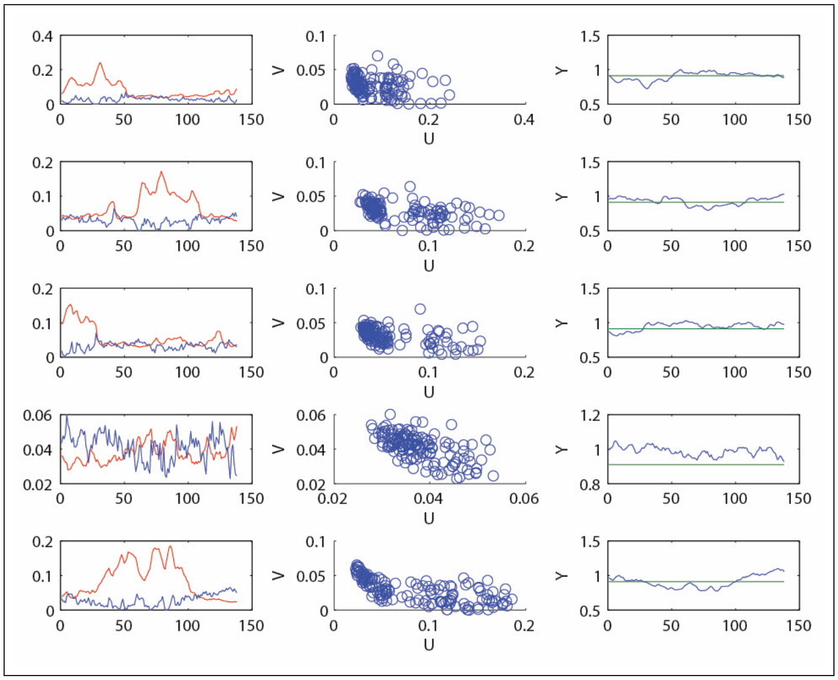

highlight shifts in the Beveridge curve that are potentially induced by the mechanism in our framework. The last column shows aggregate output Y_t for each simulation along with its threshold value for the regime switch. The graphs confirm that the simulated model reproduces the negative correlation between unemployment and vacancies; that is, the model generates a Beveridge curve. What is notable visually from the middle column of Figure 3 is that there are generally two separate clusters of data (with the exception of the sample in the fourth row). On the face of it, this lends support to the mechanism in our framework as it can replicate the shift patterns seen in actual data. This outcome is not preordained, however, as is evident from sample 4. In this simulation, the threshold is never reached despite values of output persistently below its mean for extended periods. The model economy suffers from a recession, but not one that is deep enough to do damage to the labor market.[8] Consequently, a stable Beveridge curve

[8] This is consistent with the interpretation of Benati and Lubik (2014) that most shifts of the Beveridge curve during recessions are too small to be plausibly and statis-

pattern arises over the full simulation period. We also note that this sample stands out in Table 2 because of the low volatilities of unemployment and vacancies.[9]

In the other sample paths, output falls below the threshold for lengthy periods. For instance, in the first row, the initial productivity draw pushes output below the threshold and keeps it there for fifty periods. During this period, there are two opposing forces at play. First, the productivity process is mean-reverting; that is, eventually there will be enough positive innovations to push productivity above its mean and thereby drag output back above the threshold.[10] The strength of this effect depends on the degree of persistence in productivity. If it is high enough, very large negative draws can have staying power to keep the economy below the threshold. Low persistence, on the other hand, leads to faster mean reversion. The second, endogenous force works against this pattern. If the economy is below the threshold, vacancy postings are lower than they otherwise would be (see Figure 2 and the discussion above). Consequently, matching with lower match efficiency reinforces the threshold switch. The observed Beveridge curve shift would thus be consistent with the hypothesis that prolonged periods of high unemployment are generated by mismatch in the labor market (see Şahin et al. 2014).

To summarize, we show that the simple model with an endogenous threshold switch can qualitatively and, with some qualifications, quantitatively replicate the business cycle patterns of key labor market variables. More importantly for our purposes, we demonstrate that our model can generate structural shifts in the Beveridge curve. This raises two questions. First, are these shifts large enough to be statistically different from a standard adjustment pattern, such as a counterclockwise loop as discussed in Blanchard and Diamond (1989)? Using different methodologies and sample periods, Lubik (2013) and Benati and Lubik (2014) answered this question in the negative. In this paper, we ask a second question, namely whether the shifts are even detectable as such in a flexible time-series framework.

---

tically judged as structural. They are thus more consistent with the counterclockwise loop identified by Blanchard and Diamond (1989). Yet, a few recessions, notably the most recent one, fall outside this pattern.

[9] What drives this pattern, that is, the lack of a Beveridge curve shift, is the combination of not-large-enough random shocks in the simulation and a lack of adjustment dynamics to the new (conditional) steady state associated with lower match efficiency.

[10] An alternative specification would have productivity also obey a threshold switch, so that the effect of very bad recessions would be much more protracted. This would render the model closer to the implications of a Markov-switching model such as Hamilton (1989).

A TVP-VAR for the Simulated Data

Given the simulated data, we now turn to assessing whether statistical approaches can uncover the underlying shifts in the Beveridge curve. For this purpose, we rely on a TVP-VAR with stochastic volatility, which has proved to be a flexible and useful tool to study nonlinear behavior in aggregate time series. It has recently been applied to the question of Beveridge curve shifts by Benati and Lubik (2014). Our specific time-series model builds on Cogley and Sargent (2005) and Primiceri (2005). The exposition below follows Lubik and Matthes (2015), who provide further details on the implementation.

We stack the unemployment rate U_t and the vacancy rate V_t in a column vector y_t, which we assume is determined by the following law of motion:

$$y_t = \mu_t + \sum_{j=1}^{L} A_{j,t} y_{t-j} + e_t. \tag{14}$$

μ_t is a drift term that can contain deterministic and stochastic components. The $A_{j,t}$ are conformable coefficient matrices that contain time-varying parameters. e_t is a vector of residuals. Most of the literature on TVP-VARs that use quarterly data pick the lag length in the reduced-form specification as $L = 2$. We follow this convention since we use a quarterly calibration for our matching model. We define $X_t' \equiv I \otimes (1, y_{t-1}' ..., y_{t-L}')$ to provide a concise representation of the dynamics of y_t. We thus rewrite equation (14) as:

$$y_t = X_t' \theta_t + e_t. \tag{15}$$

We assume that the law of motion for the time-varying parameters in the coefficient matrices $A_{j,t}$ is given by:

$$\theta_t = \theta_{t-1} + u_t, \tag{16}$$

where u_t is a zero mean i.i.d. Gaussian process. To characterize stochastic volatility, we assume that the covariance matrix of the one-step-ahead forecast error e_t can be decomposed using two matrices such that:

$$e_t = \Lambda_t^{-1} \Sigma_t \varepsilon_t, \tag{17}$$

where the standardized residuals are distributed as $\varepsilon_t \sim N(0, I)$. Λ_t is a lower triangular matrix with ones on the main diagonal and representative nonfixed element λ_t^i. Σ_t is a diagonal matrix with representative nonfixed element σ_t^j. The dynamics of the nonfixed elements of Λ_t and Σ_t are given by:

$$\lambda_t^i = \lambda_{t-1}^i + \zeta_t^i. \tag{18}$$

$$\log \sigma_t^j = \log \sigma_{t-1}^j + \eta_t^j. \tag{19}$$

We assume that all these innovations are normally distributed with covariance matrix V. In order to provide some structure for the estimation, we restrict the joint behavior of the innovations as follows (see Primiceri 2005):

$$V = Var\left[\begin{pmatrix} \varepsilon_t \\ u_t \\ \zeta_t \\ \eta_t \end{pmatrix}\right] = \begin{bmatrix} I & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & W \end{bmatrix}. \tag{20}$$

S is further restricted to be block diagonal, which simplifies inference. We use a Gibbs-sampling algorithm to generate draws from the posterior. The implementation of the Gibbs-sampling approach used for Bayesian inference follows Del Negro and Primiceri (2013).

A key choice for TVP-VAR modeling is how to set the prior. In order to achieve sharp inference, given the multiple sources of variation in TVP-VAR models, a researcher needs to impose restrictions on the relationship between the covariance matrices of the parameters. The trade-off, however, is that a too restrictive prior may not leave room for the time variation to appear. In our benchmark, we impose a typical choice of prior as recommended in, for instance, Primiceri (2005). Specifically, we assume the following:

$$Q \sim IW(\kappa_Q^2 * 40 * V(\theta_{OLS}), 40), \tag{21}$$

$$W \sim IW(\kappa_W^2 * 2 * I, 2), \tag{22}$$

$$S \sim IW(\kappa_S^2 * 2 * V(\Lambda_{OLS}), 2), \tag{23}$$

where IW denotes the Inverted Wishart distribution priors for all other parameters are the same as in Primiceri (2005). For the prior hyperparameters $\kappa_Q, \kappa_{W,}$ and $\kappa_{S,}$ we use the values $\kappa_Q = 0.01$, $\kappa_W = 0.01$, and $\kappa_S = 0.1$. We will discuss alternative prior choices below.

3. ESTIMATION RESULTS

We report estimation results for our benchmark TVP-VAR on simulated unemployment and vacancies data in Figures 4 and 5. In each figure, we report posterior mean estimates from the five representative data samples discussed in the previous section. Since we specify a two-variable VAR with two lags, we report eight series overall for the lag coefficients, two series for the variances, and one for the covariance.

Figure 4 Posterior Means of VAR Coefficients

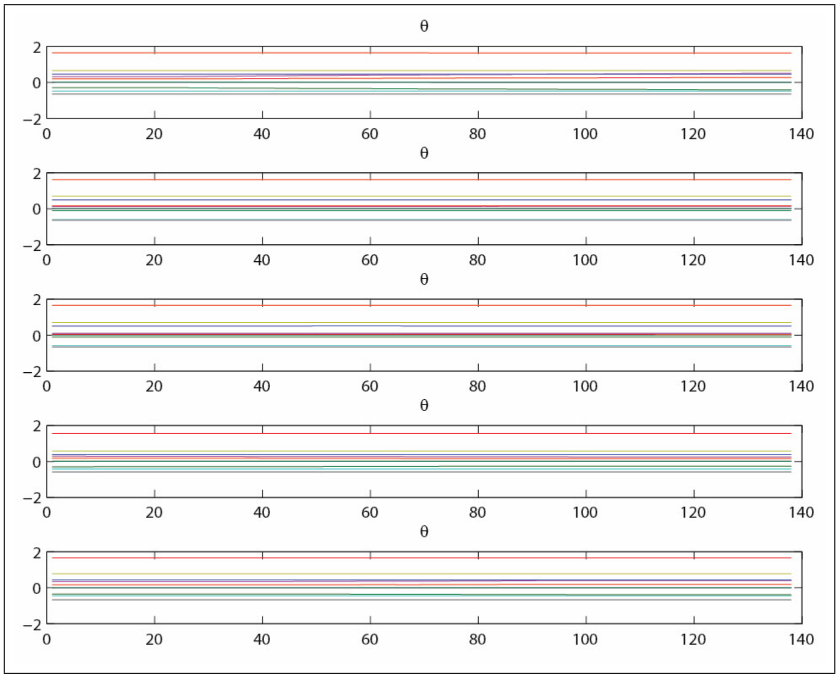

Figure 4 shows the median posterior estimates of the coefficients in the lagged matrices $A_{j,t}$ in (14) for each sample and over the entire sampling horizon. Figure 5 shows additional estimated statistics. The left column of Figure 5 reports the estimated off-diagonal elements of the covariance matrix of the one-step-ahead forecast errors, while the middle column depicts the posterior means of the diagonal elements, that is, the variances. We also report the implied regression coefficients of a period-by-period population regression of unemployment on vacancies for each sample.

Figure 5 Summary of Benchmark Results: Estimated Posterior Means

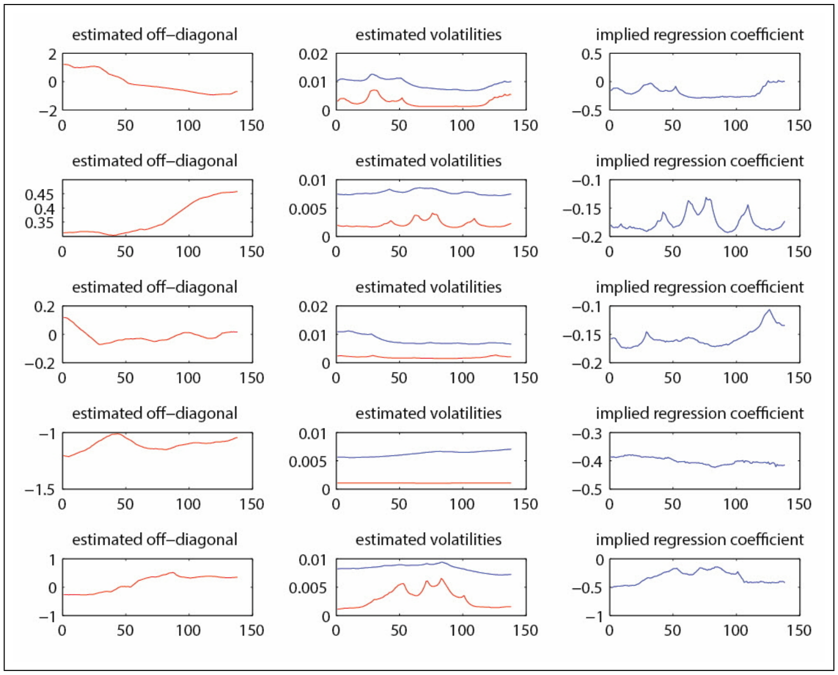

The results are almost unequivocal. Across all simulations, the TVP-VAR attributes the shifts in the simulated Beveridge curve to changes in the forecast error variance only. While both volatilities and contemporaneous correlations change with shifts in the underlying series, *all* lag coefficients are estimated to be unvarying and effectively constant (see Figure 4). The estimates for the individual samples show that when there appears to be a shift in the Beveridge curve it is associated with a gradual drift in the coefficients of the variance-covariance matrix. Consider as a baseline case the simulated sample in the fourth row of Figures 3 and 5. As discussed before, this sample path includes declines in output that never cross the threshold and therefore do not lead to Beveridge curve shifts. The TVP-VAR produces essentially constant variances of the shock innovations and an implied population regression coefficient (i.e., a Beveridge curve slope) that is fairly constant at -0.4. There is some variation in the covariance, which rises from -1.2 to -1.0 before retreating again. This seems commensurate with the increase in unemployment and the fall in vacancy postings as the economy enters a downturn in the first half of the simulated sample. The resulting pattern is that of a movement *along* the Beveridge curve but not a shift.

These patterns are noticeably different when we consider sample paths that include movements of output across the threshold. First, the population regression coefficients exhibit more variation and are smaller (in absolute value) over the full sample period compared to those of a sample path that does include a switch. Along a given Beveridge curve, unemployment and vacancies move in opposite directions. But in the transition between the two Beveridge curves, unemployment and vacancies tend to move in the same direction as vacancy postings rise in order to counteract the lower match efficiency. Shifts in the Beveridge curve are associated with shifts in the elements of the covariance matrix. In particular, periods of high volatility and positive covariation are associated with unemployment-vacancy combinations arising from low match efficiency. As discussed above, because of the constant mean and the mean-reversion of the productivity process, large and persistent enough negative shocks are required to push output below the threshold. These shocks also induce high volatility in unemployment and vacancies. The TVP-VAR then attributes this increased volatility to time-variation in the innovation covariance matrix. The positive correlation in the innovations thus mirrors the lower implied regression coefficient.

To summarize our findings, we posit that an econometrician who attempts to discover shifts in the Beveridge curve using a standard TVP-VAR would come to an erroneous conclusion. What appears in the data as a parallel shift in the curve is interpreted by the TVP-VAR as the outcome of time-variation in the variance-covariance matrix of the shocks. Large shocks drive the labor market variables away from their present location. Given the inherent persistence in the search and matching model, this would then cluster temporally close data points in a pattern that indicated a shift.[11] In the logic of the search and matching model, this outcome would be consistent with a higher incidence and severity of shocks that primarily affect the matching process and transitional labor dynamics as captured in equation (4) (see Barlevy 2011; Lubik 2013).

[11] Incidentally, this reasoning is consistent with the argument in Lubik (2013) that the degree of estimation, parameter, and model uncertainty in the empirical model is large enough that it would be difficult to distinguish statistically between competing hypotheses, especially when compared to the relatively short time span of a Beveridge curve cycle in the data. On the other hand, Benati and Lubik (2014) impose further restrictions and utilize longer samples to show that a few Beveridge curve cycles do allow for sharper inference, including the Great Recession.

However, and to reiterate this point, the underlying data are generated from a model where the presence of a structural shift for lengthy periods of time is quite noticeable. The TVP-VAR thus attributes these shifts erroneously to changes in volatility. This observation is consistent with many studies using TVP-VARs that tend not to find substantial changes in the lag coefficient matrices, but rather apportion excess volatility and breaks in behavior to stochastic volatility. Our finding is also reminiscent of the critique by Benati and Surico (2009) of Sims and Zha's (2006) argument that the switch from the Great Inflation of the 1970s to the Great Moderation of the 1980s and beyond was not driven by a break in policy but by a decline in the volatility of the shocks. By means of a simulation study, Benati and Surico (2009) show that a regime-switching VAR cannot recover a break in policy coefficients in the underlying model. Instead, it erroneously attributes the change in reduced-form behavior to changes in the innovation variance, in a manner similar to our results.

This naturally leads to the deeper question of why the TVP-VAR we use is not capable of picking up these shifts seen in the theoretical model. TVP-VARs are a very flexible modeling framework that, in theory, can certainly capture substantial shifts in parameters. At the same time, they also possess many moving parts, and the contribution of each to the ultimate estimation result is not trivial to disentangle. One important aspect is certainly the length of the sample over which the model is estimated. It is well-known that inference under heteroskedasticity (or time variation in the innovation covariance matrix) is quite problematic in short sample (e.g., Toyoda 1974). For that reason, TVP-VARs generally perform better in longer samples, as in Amir-Ahmadi, Matthes, and Wang (2016). A second aspect is that in models with many parameters, the choice of priors can be very important. In particular, priors in TVP-VARs encode a particular view of how much of the variation in the data is due to changes in parameters, changes in volatilities, or additive shocks. The following section shows one alternative to the standard practice that could be used to elicit priors. With standard priors, we would need drastic and sudden shifts in the data to have the estimated coefficients move substantially. Our search and matching model can generate those shifts, but they would arguably not be regarded as realistic for many developed economies.

4. ELICITING PRIORS FOR A TVP-VAR

A key element of TVP-VAR modeling is the choice of the prior on the time-varying components. In our benchmark specification, we follow the generally accepted practice in the literature going back to Primiceri

(2005). However, for data with considerably different properties than those commonly used in the literature or to which TVP-VAR models have been applied, our chosen values might not capture a researcher's prior view on time variation in the data set at hand. More specifically, the prior on the lag coefficient matrices $A_{j,t}$ may be too tight in our framework, so the true underlying time variation in the reduced-form coefficients is instead forced into the covariance matrix. We therefore consider an alternative that is based on a prior predictive analysis.[12]

Our alternative approach proceeds as follows. We first estimate fixed-coefficient VARs on rolling samples of the same length as our training sample (forty periods) to get paths for the time-varying coefficients and volatilities. In a separate exercise, we then simulate paths for those coefficients based on the benchmark priors described above. The hyperparameters of the alternative prior are chosen to match a set of moments from the paths of the time-varying coefficients and volatilities obtained from the rolling window estimation. We choose the average volatilities of the three sets of time-varying coefficients and volatilities. For each set of κ values that govern the tightness of the prior distribution on the covariance matrix, we run twenty-five simulations to generate paths of the same length as the paths from our rolling window estimation and average over the moments obtained in those simulations. We then pick the vector of κ coefficients that minimizes the quadratic distance between the moments from the simulations and the rolling window estimation. The difference in the moments obtained by simulation and the rolling window estimation is rescaled by their value obtained in the rolling window estimation. This avoids one set of moments dominating our calculation since the coefficients have different scales. We use a grid of values for the κ parameters. As lower bounds for the grid, we impose the values used by Primiceri (2005) since we are worried about not capturing enough time variation. Upper bounds are roughly ten times the values chosen by Primiceri (2005).

Figure 6 shows the resulting values for the prior hyperparameters for our full set of twenty-five simulated samples. The horizontal green line shows our benchmark values. As it turns out, there are, in fact, substantial differences between the values chosen by Primiceri (2005) and the values implied by our approach. In the case of the innovation variance in the law of motion for the time-varying parameters in the coefficient matrices $A_{j,t}$, equation (16), there are only two samples for which our prior choice deviates from the one chosen by this approach;

[12] An alternative would be to directly estimate the prior hyperparameters with the rest of the parameters of the model. A Gibbs sampler to do this is described in Amir-Ahmadi, Matthes, and Wang (2017).

Figure 6 Eliciting Priors: Values of Prior Hyperparameters

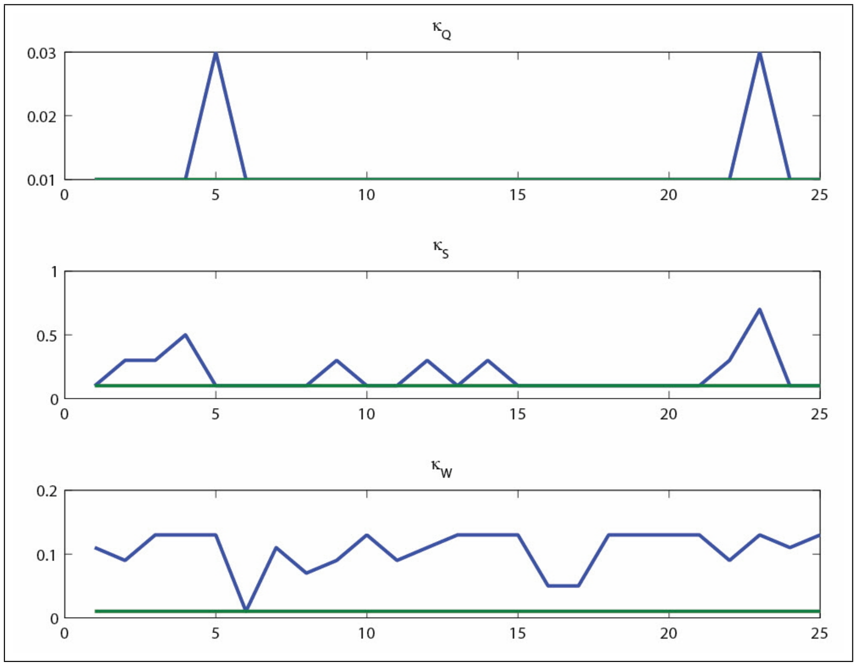

the deviation is only $\kappa_Q = 0.03$ when compared to our benchmark choice of $\kappa_Q = 0.01$. The hyperparameter scaling the variance of the innovation in the triangular decomposition matrix of the forecast-errors covariance matrix shows more deviations. They can be larger by a factor of up to five, but this is not consistent across each simulation. The largest difference to our benchmark choice can be found for the innovations on the process of the error variances. As can be seen from the bottom graph in Figure 6, the hyperparameter κ_W is larger by an order of magnitude.

This raises the question of whether this alternative prior choice has an effect on the implications derived from the TVP-VAR. We therefore reestimate our model with the much wider priors chosen by the procedure described above. The estimation results are reported in Figure 7, which can be compared directly with Figure 5. We find that the parameter estimates using the alternative prior are almost identical to those obtained using the benchmark specification. The estimated entries of the VAR companion form matrix (not reported) are also virtually identical to the benchmark case. Our conclusion that the Beveridge curve

Figure 7 Summary of Results: Estimated Posterior Means from Alternative Choice of Hyperparameters

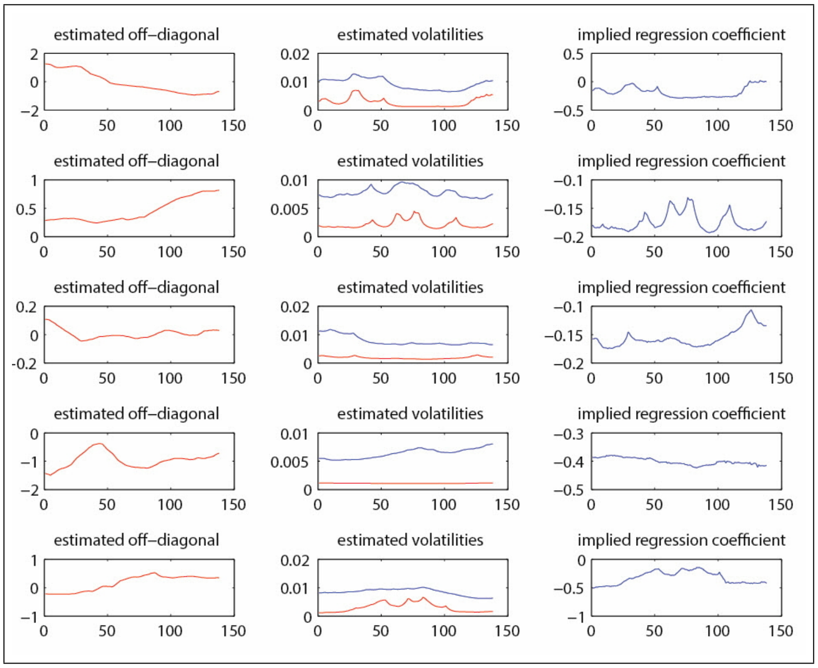

shifts in the simulated data are erroneously attributed by the TVP-VAR to time variation in the covariance matrix of the one-step-ahead prediction errors therefore remains intact. In order to get substantial differences in estimated parameters, the prior hyperparameters need to be increased dramatically (e.g. $\kappa_Q = 1$). For our application, the benchmark values consistent with the existing literature therefore seem to be a good choice as far as a naive exercise—that is, without knowledge of the underlying dynamics—is concerned.

5. CONCLUSION

This article makes a simple point. TVP-VARs appear to be predisposed to capture time variation in the underlying data by means of changes in the innovation terms and not via movements in lag coefficients. We arrived at this conclusion by means of a simulation study where we generate a specific form of nonlinearity that would imply time variation in the data. This conclusion holds for a standard choice of priors as well as an alternative set of priors that we obtain from a prior predictive analysis.

Naturally, the results derived in this article are model dependent and should therefore be taken with a grain of salt. As our model analysis shows, the degree of nonlinearity in the policy function or in the simulated data does not appear to be, heuristically speaking, large. It thus may very well be that the posterior sampler in the Bayesian estimation attributes this type of variation in the data to residual shocks, just as a fixed-coefficient VAR would. What supports this argument is that during times of economic upheaval, chiefly the Great Depression period, TVP-VARs do tend to exhibit considerable time variation in the lag coefficients (Benati and Lubik 2014; Amir-Ahmadi, Matthes, and Wang 2016). That said, we argue that the basic point still applies as to the interpretability of TVP-VAR results. At the very least, researchers should consider a more careful approach to prior selection.

In addition, and independently of the TVP-VAR angle, we propose in this article a modeling framework that conceptualizes structural changes in the labor market and links them to business cycle movements. The mechanism works via an endogenous regime shift in a key labor market parameter, whereby the shift is driven by the interaction of shocks and the intrinsic dynamics of the model. In the case of a simple labor market model, we show that a deep and long recession that originates in adverse productivity realizations can be prolonged by deterioration in the labor market matching process. This mechanism thus offers a convenient setup for studying the behavior of the labor market over the business cycle.

REFERENCES

Amir-Ahmadi, Pooyan, Christian Matthes, and Mu-Chun Wang. 2016. "Drifts and Volatilities under Measurement Error: Assessing Monetary Policy Shocks over the Last Century." *Quantitative Economics* 7 (July): 591–611.

Amir-Ahmadi, Pooyan, Christian Matthes, and Mu-Chun Wang. 2017. "Estimating Prior Hyperparameters in VAR Models with Time-Varying Parameters and Stochastic Volatility." Manuscript.

Barlevy, Gadi. 2011. "Evaluating the Role of Labor Market Mismatch in Rising Unemployment." Federal Reserve Bank of Chicago *Economic Perspectives* 35 (Third Quarter): 82–96.

Benati, Luca, and Thomas A. Lubik. 2014. "The Time-Varying Beveridge Curve." In *Advances in Non-linear Economic Modeling: Theory and Applications*, edited by Frauke Schleer-van Gellecom. Heidelberg: Springer Berlin Heidelberg, 167–204.

Benati, Luca, and Paolo Surico. 2009. "VAR Analysis and the Great Moderation." *American Economic Review* 99 (September): 1636–52.

Blanchard, Olivier J., and Peter A. Diamond. 1989. "The Beveridge Curve." *Brookings Papers on Economic Activity* (No. 1): 1–76.

Canova, Fabio, Filippo Ferroni, and Christian Matthes. 2015. "Approximating Time Varying Structural Models With Time Invariant Structures." Federal Reserve Bank of Richmond Working Paper 15-10 (September).

Cogley, Timothy, and Thomas J. Sargent. 2005. "Drift and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S." *Review of Economic Dynamics* 8 (April): 262–302.

Del Negro, Marco, and Giorgio Primiceri. 2013. "Time-Varying Structural Vector Autoregressions and Monetary Policy: A Corrigendum." Federal Reserve Bank of New York Staff Reports 619 (May).

Hagedorn, Marcus, and Iourii Manovskii. 2008. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited." *American Economic Review* 98 (September): 1692–706.

Hamilton, James D. 1989. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica* 57 (March): 357-84.

Lubik, Thomas A. 2009. "Estimating a Search and Matching Model of the Aggregate Labor Market." Federal Reserve Bank of Richmond *Economic Quarterly* 95 (Spring): 101–20.

Lubik, Thomas A. 2013. "The Shifting and Twisting Beveridge Curve: An Aggregate Perspective." Federal Reserve Bank of Richmond Working Paper 13-16 (October).

Lubik, Thomas A., and Christian Matthes. 2015. "Time-Varying Parameter Vector Autoregressions: Specification, Estimation, and an Application." Federal Reserve Bank of Richmond *Economic Quarterly* 101 (Fourth Quarter): 323–52.

Primiceri, Giorgio. 2005. "Time Varying Structural Vector Autoregressions and Monetary Policy." *Review of Economic Studies* 72 (July): 821–52.

Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L. Violante. 2014. "Mismatch Unemployment." *American Economic Review* 104 (November): 3529–64.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95 (March): 25–49.

Sims, Christopher A., and Tao Zha. 2006. "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review* 96 (March): 54–81.

Toyoda, Toshihisa. 1974. "Use of the Chow Test Under Heteroscedasticity." *Econometrica* 42 (May): 601–8.

APPENDIX: DERIVATION OF THE WAGE SCHEDULE

The wage that firms pay to workers is derived as the outcome of a Nash bargaining process. Denoting the workers' weight in the bargaining process as $\eta \in [0, 1]$, this implies the sharing rule:

$$\mathcal{W}_t - \mathcal{U}_t = \frac{\eta}{1-\eta}\mathcal{J}_t, \tag{A1}$$

where $\mathcal{W}_t$ is the asset value of employment, $\mathcal{U}_t$ is the value of being unemployed, and $\mathcal{J}_t$ is, as before, the value of the marginal worker to the firm. In models with one-worker firms, the net surplus of a firm is given by $\mathcal{J}_t - \mathcal{V}_t$, with $\mathcal{V}_t$ the value of a vacant job. By free entry, $\mathcal{V}_t$ is then assumed to be driven to zero. The value of employment to a worker is described by the following Bellman equation:

$$\mathcal{W}_t = w_t + \beta E_t[(1 - \chi_{t+1})\mathcal{W}_{t+1} + \chi_{t+1}\mathcal{U}_{t+1}]. \tag{A2}$$

Workers receive the wage w_t and transition into unemployment in the next period with probability s. The value of searching for a job, when currently unemployed, is:

$$\mathcal{U}_t = b + \beta E_t[\chi_t(1 - \chi_{t+1})\mathcal{W}_{t+1} + (1 - \chi_t(1 - \chi_{t+1}))\mathcal{U}_{t+1}]. \tag{A3}$$

An unemployed searcher receives benefits b and transitions into employment with probability $\chi_t(1 - \chi_{t+1})$. It is adjusted for the probability that a completed match gets dissolved before production begins next period. Substituting the asset equations into the sharing rule (A1), results, after some algebra, in the wage equation found in the text:

$$W_t = \eta\,(A_t + \kappa\theta_t) + (1 - \eta)b. \tag{A4}$$

APPENDIX: MODEL SOLUTION

We solve the simple search and matching model fully nonlinearly by means of the monotone mapping algorithm. The algorithm computes an approximation of the decision rule $\widehat{h}^V(N_t, A_t, \chi_t)$, which determines the number of vacancy postings given the state variables: employment N_t, the exogenous productivity shock A_t, and the separation rate process χ_t. The algorithm contains the following steps:

1. Specify a threshold switching value $\underline{Y}$ and discretize the state space S. Formulate an initial guess for the decision rule: $\widehat{h}_0^V(N, A, \chi)$ $\forall\{N, A, \chi\} \in S$.
2. Compute a residual function $\overline{R}(V_t; \{N_t, A_t, \chi_t\})$ based on the following:
   - (a) Y_t is calculated and m_t is given according to the threshold process (1).
   - (b) Calculate next period's employment from (4).
   - (c) Expected values of next-period values in the firm's first-order condition appear as:

$$E_t\left[\frac{(1-\eta)\left(A_t^{\rho_A}e^{\varepsilon_{At}}-b\right)-\eta\kappa(\frac{V_{t+1}}{1-N_{t+1}})+(\frac{\kappa}{m_{t+1}})(\frac{V_{t+1}}{1-N_{t+1}})^{\xi}}{Y_{t+1}}\right] = X_t,$$

which can be approximated with the truncated distribution:

$$\widehat{X}_t = \int_{\underline{\varepsilon}}^{\varepsilon^*}\phi(\varepsilon;\sigma_\varepsilon^2)\frac{\Phi\left(N_{t+1}, A_t\right)}{\Psi\left(N_{t+1}, A_t\right)}d\varepsilon + \int_{\varepsilon^*}^{\overline{\varepsilon}}\phi(\varepsilon;\sigma_\varepsilon^2)\frac{\Phi\left(N_{t+1}, A_t\right)}{\Psi\left(N_{t+1}, A_t\right)}d\varepsilon,$$

where:

$$\begin{aligned}\Phi\left(N_{t+1}, A_t\right) &= (1-\eta)\left(A_t^{\rho_A}e^{\varepsilon_{At}}-b\right) - \\ &\quad \eta\kappa(\frac{\widehat{h}_0^V(N_{t+1}, A_t^{\rho_A}e^{\varepsilon_{At}})}{1-N_{t+1}}) + \\ &\quad +(\frac{\kappa}{m_l})(\frac{\widehat{h}_0^V(N_{t+1}, A_t^{\rho_A}e^{\varepsilon_{At}})}{1-N_{t+1}})^{\xi}, \\ \Psi\left(N_{t+1}, A_t\right) &= A_t^{\rho_A}e^{\varepsilon_{At}}N_{t+1} - \kappa\widehat{h}_0^V(N_{t+1}, A_t^{\rho_A}e^{\varepsilon_{At}}).\end{aligned}$$

Estimate this expression with a trapezoid rule. Linear interpolation is used in the implementation of the decision rule.

(d) Given the expectations, the residual function is:

$$\overline{R}(V_t; \{N_t, A_t, \chi_t\}) = \left|(1-\chi_t)\beta Y_t\ \widehat{X} - \frac{\kappa}{m_t}\left(\frac{V_t}{1-N_t}\right)^{\xi}\right|,$$

which can be interpreted as the absolute value of the difference between the right-hand side and left-hand side of the firm's first-order condition.

3. This residual function is minimized over V_t for every triple $\{N_i, A_j, \chi_k\}$ in S. The decision rule is then updated based on:

$$\widehat{h}_2^V(N_i, A_j, \chi_k) = \arg\min\{\overline{R}(V_t; \{N_i, A_j, \chi_k\})\}\ \forall\{n_i, A_j, \chi_k\} \in S.$$

4. The algorithm is repeated until:

$$\max\left|\widehat{h}_{k+1}^V(N_i, A_j, \chi_k) - \widehat{h}_k^V(N_i, A_j, \chi_k)\right| < \varepsilon.$$

Economic Quarterly—Volume 102, Number 3—Third Quarter 2016—Pages 227–252

The Heterogeneous Business-Cycle Behavior of Industrial Production

Jackson Evert and Felipe Schwartzman

Industry-level data can provide a window into the sources of business cycles as well as propagation mechanisms. This is because depending on what determines those, one might expect different industries to behave differently. One notable example of the use of industry-level data for that purpose is Gertler and Gilchrist (1994), who pointed to the relatively larger impact of monetary shocks in industries with relatively smaller sized firms as evidence for the role of financial frictions in propagating those shocks. Another example is Bils et al.'s (2013) comparison of markup fluctuations in durable vs. nondurable sectors as a means to assess whether demand fluctuations could cause fluctuations in markups.

The use of industry-level variation can also provide advantages over the use of even more disaggregated firm-level data. First, since industries are to a large extent defined by the nature of their products, differences between industries are more plausibly determined by stable differences in technology and preferences than differences across firms within an industry. Second, because industry-level data already allow for some aggregation, they capture at least part of the general equilibrium effects that are likely to be important at the aggregate level. Third, industry-level data are more readily available, allowing for a useful first pass before acquiring harder-to-obtain firm-level data. The clear disadvantage is that because industries are different from one

We thank Bob Hetzel, Christian Matthes, Daniel Schwam, and John Weinberg for helpful comments and suggestions. The views expressed in this article are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Richmond or the Federal Reserve System. All errors are our own. E-mail: felipe.schwartzman@rich.frb.org.

DOI: http://doi.org/10.21144/eq1020303

another along several dimensions, one needs to be always concerned about the possibility that industry variation is driven by some omitted characteristic. Thus, any work using industry-level data must incorporate extensive controls.

The purpose of this article is to present some stylized facts for how the business-cycle behavior of sectoral output differs with sectoral characteristics. Those stylized facts can be informative either as a means to determine sources of fluctuations and transmission channels or as indications of important sources of sectoral heterogeneity that ought to be controlled for in any study that attempts to uncover those sources and channels. We construct these stylized facts by first calculating standard business-cycle statistics such as relative volatility and correlation with GDP for each of the seventy-two sectors for which industrial production data are available separately. With those statistics in hand, we can then ask which industry-level characteristics are most likely to predict how these moments vary.

The measures of sectoral characteristics we focus on fall into four categories. The first category includes determinants of the demand for products in different sectors. Those may be informative about the role of fluctuations in the composition of demand for different types of products on business cycles. For example, the extent to which sectors that have the government as a main customer fluctuate more or less with aggregate GDP provides some information about the role of government consumption in business cycles (Ramey [2011] provides a recent review of the literature). The second category includes determinants of production costs. Those can provide a window into the role of cost fluctuations in business cycles. For example, a wide literature has pointed to energy cost fluctuations as an important driver of business cycles (see Hamilton [2003] for a seminal example). Variables in the two categories, demand and cost, can provide information about the role of the integration of different industries in production chains. This can help shed light on theories of business-cycle propagation that emphasize the input-output structure of the economy, such as Acemoglu et al. (2012). The third category includes measures of pricing distortions, including measures of market power and of price stickiness. Those can shed light on theories of business cycles that emphasize markup fluctuations as a key propagation mechanism (Rotemberg and Woodford [1999] provide a review). The fourth category includes firm-level characteristics that the literature has pointed to as correlated with sensitivity to financial frictions. Those are relevant for theories of business cycles that emphasize financial shocks and financial frictions (Bernanke and Gertler 1989; and Kiyotaki and Moore 1997). Those different categories are constructed in order to obtain a wide scope of cross-industry differences that the existing literature has pointed out as potentially important.

Some of the most salient findings are as follows:

1) Industries that are more oriented toward the production of consumer goods, which produce goods that are nondurable, and that produce necessities tend to be less volatile and less correlated with business cycles than other industries. Furthermore, they also tend to lead them. A similar pattern is present in firms that intensively use agricultural inputs.

2) Industries that are more oriented toward the production of goods consumed by the government are less correlated with business cycles relative to other industries and tend to lag business cycles. At the same time, industries that are more oriented toward the private sector tend to lead business cycles.

3) Industries in which nominal prices change infrequently tend to lag business cycles.

4) Industries whose characteristics are likely to be correlated with sensitivity to financial frictions are likely to lag business cycles, whereas those that are less likely to be exposed to those frictions tend to lead them.

5) The position of different industries in the production chain matters. Industries that are highly integrated in the production chain either by being intensive in the use of intermediate inputs or by dedicating a large fraction of their output to intermediate inputs are more likely to lead GDP.

The first section provides a more careful description and justification of the methodology. The subsequent section represents the core of the paper. First, it presents a description of how the different moments are distributed across sectors. Then, in four subsections we provide more detail on the findings for each of the four categories described above and provide some discussion of those findings in light of existing literature. After those, we perform a multivariate analysis to account for the fact that industry characteristics might be correlated among themselves. The last section summarizes the results. In the Appendix, we present a detailed description of how we constructed the various measures of industry characteristics.

1. METHODOLOGICAL DETAILS

In this section, and in all sections that follow, we will examine statistics for detrended time series. The detrending process follows Hodrick and Prescott (1997) and involves fitting a curve through the time series that strikes a balance between staying close to the data and remaining relatively smooth.[1] This trade-off is controlled by a parameter that, in one extreme, makes the estimated trend perfectly smooth and, hence, linear and, on the other extreme, leads to an estimated trend that is identical to the data. The commonly used parameter for quarterly data is 1600. The detrended series is then the log difference between the series and the estimated trend. In what follows we refer to a moment as being a "business-cycle" moment whenever it is constructed using HP-filtered time series.

In order to gather a better understanding of how different moments provide different information about the comovement of sectoral output and business cycles, consider first the following model of detrended sectoral output in which, for simplicity, we abstract from dynamics:

$$Y_{i,t} = \sum_{r=1}^{R} \lambda_{i,r} \epsilon_{r,t},$$

where $Y_{i,t}$ is output in sector i, $\epsilon_{r,t}$ are the values at time t of each of R shocks potentially affecting all sectors, and $\lambda_{i,r}$ is the sensitivity of sectoral output to each of the aggregate shocks. Shocks $\epsilon_{r,t}$ are uncorrelated with one another, i.e., $cov(\epsilon_{r,t}, \epsilon_{r',t}) = 0$ for all $r \neq r'$ and all t. Note that this specification is quite flexible, since we do not restrict R to be a small number relative to the number of sectors. In particular, the shocks $\epsilon_{r,t}$ can include idiosyncratic shocks, i.e., shocks that affect only one sector. It also accommodates setups in which shocks that affect primarily one sector also affect other sectors through input-output linkages, etc.[2] For simplicity, assume that detrended aggregate output can be approximated as a simple average of sectoral output, so that

$$Y_t = \sum_{i=1}^{I} \frac{Y_{i,t}}{I}.$$

The simplest moment of interest is the business-cycle variance of sectoral output relative to that of aggregate GDP. If we normalize the variance of the aggregate shocks $\epsilon_{r,t}$ to 1, this is

[1] As a robustness test, we also generated the tables using a Band-Pass filter (see Baxter and King [1999] for details on that kind of filtering). They are available upon request.

[2] See Acemoglu et al. (2012) for analytical and quantitative explorations. We refer the reader to these papers for further details. For the purposes of this essay, one can accommodate that view by reinterpreting some of the aggregate shocks as shocks that affect primarily particular sectors but do not "wash out" in aggregate due to linkages.

$$\frac{std(Y_{i,t})}{std(Y_t)} = \sqrt{\frac{\sum_{r=1}^{R} \lambda_{i,r}^2}{\sum_{r=1}^{R} (\sum_{i=1}^{I} \lambda_{i,r}/I)^2}}$$

or, more compactly,

$$\frac{std(Y_{i,t})}{std(Y_t)} = \sqrt{\frac{\sum_{r=1}^{R} \lambda_{i,r}^2}{\sum_{r=1}^{R} \bar{\lambda}_r^2}},$$

where $\bar{\lambda}_r \equiv \sum_{i=1}^{I} \lambda_{i,r}/I$ is the average sensitivity of sector i to aggregate shock r. In this benchmark case, the relative variance of a sector is large if $\lambda_{i,r}^2$ is on average large relative to $\bar{\lambda}_r^2$. Note that this measure does not allow us to distinguish whether the large relative variance stems from a relatively large sensitivity to shocks that are also important for other sectors (i.e., $\lambda_{i,r} > \bar{\lambda}_r >> 0$) or from a high sensitivity to a shock that is not relevant for other sectors (i.e., $\lambda_{i,r} > \bar{\lambda}_r \simeq 0$). The latter case would correspond to a case in which sector-specific shocks are very large for individual sectors as compared to aggregate shocks but “wash-out” in aggregate.

The correlation of industrial output with GDP provides an alternative view on the cyclical sensitivity of a sector. If business cycles were predominantly caused by a single common shock to all sectors, with sector-specific shocks playing a very small role, one would expect the correlation of all sectoral output with aggregate GDP to be very close to one. Contrariwise, if sectoral shocks play a disproportionate role in individual sector output, one would expect the correlation of that sector with GDP to be relatively smaller. Similarly, one may find small correlations if output in a given sector is driven by an aggregate shock that is not the main driving force of aggregate business cycles. In terms of our simple model with $I \to \infty$, the correlation between any given sector and aggregate output is

$$corr(Y_{i,t}, Y_t) = \frac{\sum_r \lambda_{i,r}\bar{\lambda}_r}{(\sum_r \lambda_{i,r})^2(\sum_r \bar{\lambda}_r)^2}.$$

If $\lambda_{i,r}$ and $\bar{\lambda}_r$ have mean zero, the correlation between $Y_{i,t}$ and Y_t would be simply given by the correlation between between $\lambda_{i,r}$ and $\bar{\lambda}_r$. More generally, it is an increasing function of that correlation. Thus, the correlation between sectoral output and aggregate output measures the extent to which the two are driven by the same shocks.

Note that it is possible for the output of a given industry to be at the same time much more volatile than aggregate output and to have a low contemporaneous correlation. This would happen if such an industry's output is largely determined by idiosyncratic shocks, which have little effect on the output of other industries. Conversely, an industry might be less volatile than aggregate output but also highly correlated if it is mostly driven by the same shock that drives other industries but is comparatively less sensitive to those.

Finally, apart from relative variances and correlation with GDP, we also provide statistics for the correlation of sectoral output and leads and lags of output. Interpreting those requires a dynamic model. This is a straightforward generalization of the model described above, in which industrial output depends on shocks that occurred in the past:

$$Y_{i,t} = \sum_{s=0}^{\infty} \sum_{r=1}^{R} \lambda_{i,r,s} \epsilon_{r,t-s},$$

where we now also impose that $cov(\epsilon_{i,t}, \epsilon_{j,t-s}) = 0 \ \forall i, j, s$; that is, we impose that shocks are i.i.d., with all persistence a function of $\lambda_{i,r,s}$. The model above is fairly general, as it corresponds to a moving average representation of a vector-valued time-series model (see, for example, Hamilton [1994] for a detailed discussion).

Note that under this more general framework, it is possible for two variables to be contemporaneously uncorrelated even if they are driven by the same shock, so long as that occurs at different lags. For example, if $Y_{i,t} = \epsilon_{1,t}$ and $Y_{i*,t} = \epsilon_{1,t-1}$, those two processes will have zero contemporaneous correlation. However, the correlation of $Y_{i,t}$ and $Y_{i*,t+1}$ will be equal to one. More generally, examining lead and lagged correlations may provide us with some indication of whether certain industries are more likely to respond more sluggishly with shocks than overall GDP, a fact that is likely to be reflected in relatively low contemporaneous correlations by relatively high correlations with lagged output. Conversely, examining correlations with leads and lags of output may provide us a sense of variables that react more rapidly to shocks, thus forecasting output.

2. THE CROSS-SECTORAL DISTRIBUTION OF BUSINESS-CYCLE MOMENTS

Table 1 shows some descriptive statistics for the distribution of various business-cycle moments across sectors. The first observation is that in all sectors, business-cycle variance is larger than that of aggregate output, and for the median sector it is four times as large. This observation is consistent with the notion that output in individual sectors is largely

Table 1 Summary Statistics

	Variance	Mean	Median	25th Percentile	75th Percentile
Std. Dev.	3.38	3.93	4.03	2.54	4.80
t-8	0.03	-0.21	-0.23	-0.31	-0.12
t-6	0.03	-0.07	-0.08	-0.21	0.04
t-4	0.04	0.15	0.14	0.02	0.27
t-3	0.04	0.28	0.26	0.13	0.45
t-2	0.05	0.40	0.40	0.23	0.57
t-1	0.06	0.49	0.54	0.33	0.69
t	0.07	0.53	0.63	0.32	0.74
t+1	0.06	0.47	0.53	0.30	0.69
t+2	0.05	0.36	0.39	0.18	0.54
t+3	0.04	0.24	0.25	0.08	0.41
t+4	0.04	0.14	0.14	0.00	0.28
t+6	0.03	0.00	0.00	-0.16	0.15
t+8	0.03	-0.11	-0.13	-0.24	0.02

Note: The cells refer to descriptive statistics of moments across industries. For each industry, we calculate a standard deviation and correlations with leads and lags of output. We then report statistics summarizing the cross-industry distribution of those moments.

driven by idiosyncratic shocks that are to a large extent averaged out in aggregate.

The second observation is that the correlation of sectoral output with aggregate GDP is mostly positive (animal food manufacturing and dairy product manufacturing being the only sectors with a negative correlation). The median sector has a correlation of 0.63 with GDP, and 75 percent of the sectors have a correlation of more than 0.32.

Third, the median correlation with leads and lags of GDP declines as the number of leads or lags increase in a fairly symmetrical fashion. At six-quarter leads and lags, the median sector has a correlation with output that is fairly close to zero. In the next subsection, we will describe how those business-cycle moments correlate with various measures of industry characteristics.

Demand

We start our investigation of stylized facts by examining how business-cycle moments depend on determinants of sectoral demand. There is no a priori reason why the demand for different products should vary in the same way with business cycles. In fact, sectoral variation in sensitivity to different demand components can provide a way to test theories of propagation of demand shocks. For example, Bils et al.

(2013) use cross-industry variation in sensitivity of demand as a means to assess the ability of demand shocks to lead to markup variations.

It is a well-known stylized fact of business cycles that consumption of nondurable goods varies less than output over the business cycle, whereas the demand for durable consumer goods and investment goods varies more than output. This suggests that sectors whose production is more dedicated to consumption ought to experience relatively lower business-cycle variation. We check whether this simple prediction is true by constructing for each sector a measure of the importance of household consumption in its output. Roughly speaking, it corresponds to the share of the output of each industry that is purchased by households as consumer goods (see Appendix for a detailed discussion of how this and other measures are constructed). As we can observe, the prediction is born out by the data, with consumption-oriented sectors exhibiting lower business-cycle variance, although the negative correlation is relatively small in absolute value. Interestingly, however, the *correlation* of sectoral output with the business cycle also declines with its orientation toward household consumption. This suggests that compared to other sectors, sectors oriented toward household consumption are more likely to be driven by shocks other than the ones determining overall GDP. Interestingly, the pattern disappears and, in fact, reverses itself once one compares business-cycle fluctuations at the sectoral level with that of future GDP. It appears that, relatively speaking, household consumption-oriented sectors tend to *lead* business cycles. This may imply some ability on the part of households to forecast business cycle shocks and adjust their consumption accordingly early on.

Bils et al. (2013) focus on durability of the goods produced in different sectors as a major source of variation in sensitivity to demand shocks. Demand for durable goods is particularly sensitive to shocks because stocks of durables are much larger than purchases in any given period, so large changes in those purchases are necessary in order to change the stock in use. More concretely, suppose a car depreciates at a rate δ, and aggregate household demand for cars is given by $X_{car,t}$. For simplicity of exposition, suppose demand follows an exogenous process. Then, if demand for cars increases by 1 percent, this requires increasing the stock of cars in circulation by 1 percent. However, if we take a stable demand for cars as a baseline, households must increase their purchase of cars from $\delta X_{car,t}$ (the amount that they need to purchase in order to make up for depreciation) to $(\delta + 0.01)X_{car,t}$, an increase of $1/\delta$ percent. Thus, if cars depreciate at a rate of 5 percent per quarter, this implies an increase in car purchases of 20 percent. Consistently with those calculations, output volatility does seem to be tightly linked to the durability of the good produced in a given sector.

Figure 1 Demand Correlates

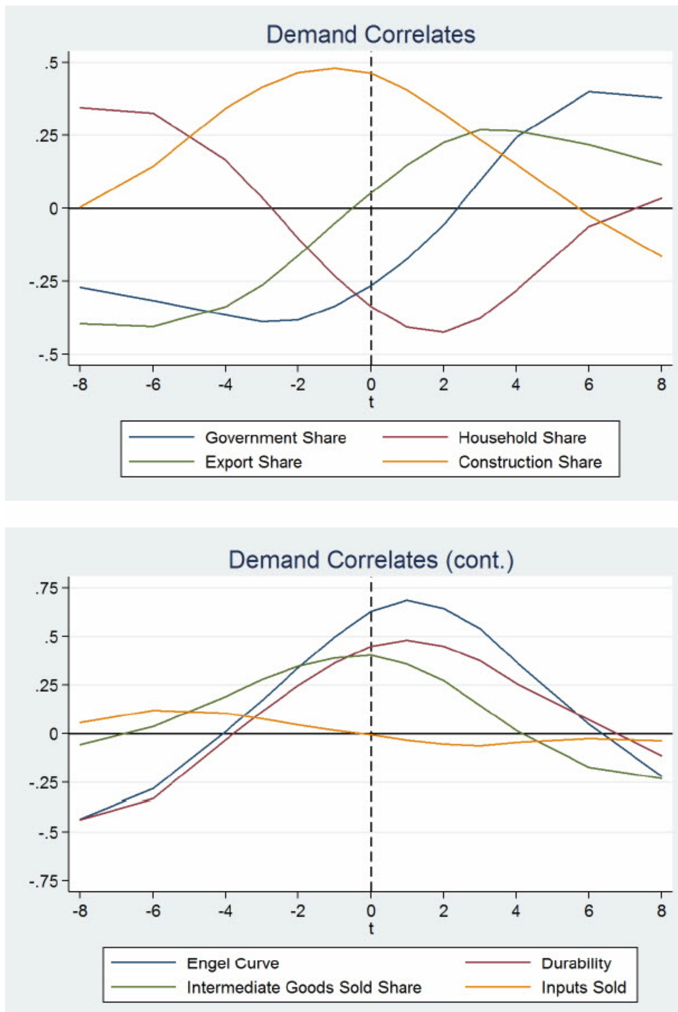

Table 2 Demand Correlates

	Std. Dev.	t-8	t-4	t-3	t-2	t-1	t	t+1	t+2	t+3	t+4	t+8
Household Share	-0.19	0.35	0.16	0.04	-0.1	-0.23	-0.34	-0.41	-0.43	-0.38	-0.28	0.03
Government Share	-0.04	-0.27	-0.37	-0.39	-0.38	-0.34	-0.26	-0.18	-0.06	0.09	0.24	0.38
Construction Share	0.02	0.00	0.34	0.41	0.46	0.48	0.46	0.4	0.32	0.23	0.15	-0.17
Export Share	0.36	-0.4	-0.34	-0.26	-0.16	-0.05	0.05	0.15	0.22	0.27	0.26	0.15
Intermediate Share	-0.15	-0.06	0.19	0.28	0.34	0.39	0.40	0.36	0.27	0.14	0.02	-0.23
Inputs Sold	-0.27	0.05	0.1	0.08	0.04	0.02	-0.01	0.04	-0.05	-0.06	-0.04	-0.04
Engel Curve	0.45	-0.44	0.01	0.17	0.34	0.5	0.63	0.68	0.64	0.54	0.37	-0.22
Durability	0.62	-0.44	-0.03	0.11	0.25	0.36	0.45	0.48	0.45	0.38	0.26	-0.12

Note: Table reports the correlations between industry characteristics and business-cycle moments (either relative volatility or business-cycle correlation for various industry leads/lags).

The findings for the correlation between depreciation and various moments largely resemble those for household consumption orientation, with sectors producing more durable goods being more contemporaneously correlated with GDP and less-durable sectors leading aggregate GDP. The main difference between the two measures is that durability is a much better predictor of the relative volatility of different sectors than consumption orientation.

Another household-demand-related dimension that one might expect to be predictive of the sensitivity of output in different sectors to business-cycle variations is the income elasticity of demand for that good (or the slope of the Engel Curve). Bils et al. (2013) estimate this elasticity using cross-sectional data. Using their estimates, we find that sectors with steeper Engel Curves are also more volatile and more correlated with output. The result is interesting in that it suggests that business-cycle variation in national income has a qualitatively similar impact on household demand composition as variation in income across households at a given point in time. It is also noteworthy that necessary goods (i.e., those with low income elasticity) are particularly good predictors of business cycles. Those goods also tend to be more household-oriented and have higher depreciation rates. Interestingly, the magnitude of the correlations between Engel coefficients and output correlations stands out when compared to the other metrics.

Given the focus of much of business-cycle analysis on the role of fiscal shocks, one further demand-side related metric of interest is orientation of a given sector toward government consumption. That metric is especially interesting since it provides a window into the role of fiscal shocks in driving sectoral output. Sectors oriented toward government consumption do not appear to be more or less volatile than other sectors. However, they are less contemporaneously correlated with business cycles, as one would expect if government purchasing decisions were largely disconnected from broader economic conditions. Interestingly, however, they become more correlated with lags, implying that the impact of shocks affecting output in most sectors only affect those that are oriented toward government consumption with delay.

The orientation of individual industries toward construction provides a further dimension of industry demand that is likely to be informative about theories of the business cycle. We find that those sectors do tend to be more correlated with business cycles, in line with theories that have gained prominence after the Great Recession, consistent with housing demand playing a prominent role in driving business-cycle fluctuations. Furthermore, they appear to lead business cycles slightly.

A further source of industry-level variation is motivated by recent work on the interplay between industry-level and aggregate dynamics, which has emphasized the importance of input-output linkages in the propagation of shocks. This suggests that it could be interesting to investigate whether industry-level business-cycle moments correlate with a measure of how "upstream" an industry is, meaning what fraction of its output is sold as inputs to other industries. We find that such industries, while not more or less volatile than others, tend to be more correlated with business cycles. They also are slightly more correlated with future output than with past output, hinting at timing delays between the production of intermediate inputs and final outputs.[3]

Finally, we investigate the extent to which the foreign orientation of a sector makes it more or less correlated with business cycles. We find that sectors that are less export-oriented tend to lead the business cycle relative to sectors that are more export-oriented. Thus, export-oriented sectors appear to be more insulated from business-cycle shocks in early stages.

Cost

We now turn to measures capturing the intensity of use of different inputs in production. We start by focusing on those input categories that are likely to have the most volatile prices, including energy, food, and mining. To the extent that industries that are intensive in those inputs are correlated with business cycles, this may indicate that shocks to the supply of these inputs may help drive business-cycle fluctuations. Of those three, the one that appears to have the most predictive power over industry-level business-cycle statistics is the fraction of agricultural inputs used in production. However, rather than implying that

[3] Following Acemoglu et al. (2012), we also examine the role, if any, of heterogeneity in industry "degree," as measured by the fraction of industry intermediate input production in total production of intermediate inputs in the economy. For that measure, we did not find that this has any predictive impact on business-cycle moments.

Figure 2 Cost Correlates

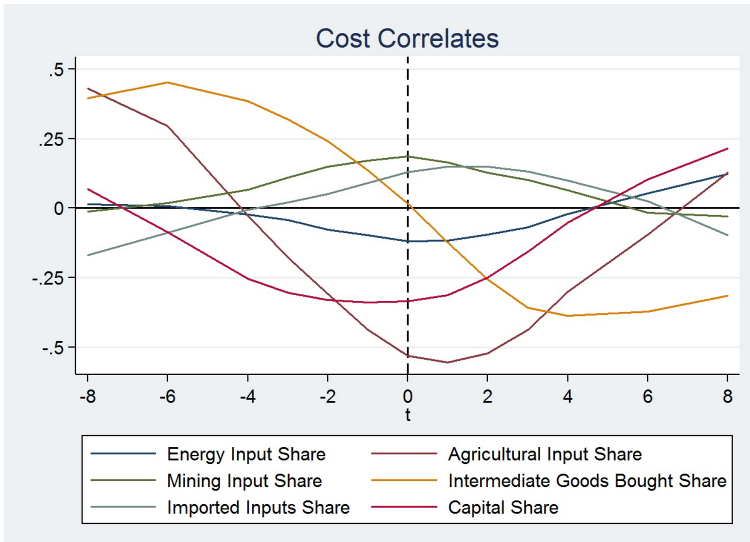

Note: Figures report the correlations between industry characteristics and business-cycle moments (either relative volatility or business-cycle correlation for various industry leads/lags).

agricultural cost shocks drive business cycles, the main finding is that industries intensive in agricultural inputs appear to be more *disconnected* from business cycles, with contemporaneous correlations being smaller the more agricultural inputs are used. Interestingly, however, their volatility is also relatively smaller. Industries with agricultural inputs also tend to lead business cycles, in a pattern reminiscent of low Engel elasticity sectors. This occurs in part because sectors that use agricultural goods in production are in part producing exactly such necessities. The multivariate analysis in Section 2.5 should help us disentangle those effects.

Table 3 Cost Correlates

	Std. Dev.	t-8	t-4	t-3	t-2	t-1	t	t+1	t+2	t+3	t+4	t+8
Energy Inputs	-0.24	0.01	-0.02	-0.04	-0.08	-0.1	-0.12	-0.12	-0.1	-0.07	-0.02	0.12
Agricultural Inputs	-0.36	0.43	-0.03	-0.18	-0.31	-0.44	-0.53	-0.55	-0.52	-0.44	-0.3	0.13
Mining Inputs	0.04	-0.01	0.07	0.11	0.15	0.17	0.19	0.16	0.13	0.1	0.06	-0.03
Intermediate Inputs	0.00	0.39	0.38	0.32	0.24	0.14	0.02	-0.12	-0.26	-0.36	-0.39	-0.32
Imported Inputs	0.30	-0.01	0.15	0.15	0.16	0.16	0.16	0.13	0.07	0.01	-0.04	-0.22
Imp. Share of Inputs	0.32	-0.17	-0.01	0.02	0.05	0.09	0.13	0.15	0.15	0.13	0.1	-0.1
Capital Share	-0.18	0.07	-0.25	-0.3	-0.33	-0.34	-0.33	-0.31	-0.25	-0.16	-0.05	0.21

Comparatively speaking, sectors with high intensity in energy and mining inputs do not seem to be more or less correlated with business cycles than other sectors. The low correlation with energy intensity is somewhat surprising in light of the common notion that energy shocks are an important source of business-cycle fluctuations. We also examine what happens when we eliminate the three industries with the highest use of energy inputs, since those have a level of energy use that is much higher than the others and are themselves involved in energy production. Eliminating those sectors does not increase the extent to which business-cycle correlations are associated with energy use.[4]

We also investigate whether capital intensity and intermediate input intensity are predictive of business-cycle correlations. Capital-intensive sectors appear to be less correlated with business cycles contemporaneously but more correlated after eight quarters. This suggests a sluggish response of those sectors to business-cycle shocks in line with capital-adjustment costs and planning lags.

Furthermore, we examine the correlation between the fraction of intermediate inputs in total output and business cycles. We find that sectors that use more intermediate inputs are no more or less correlated with business cycles than sectors that use fewer intermediate inputs. However, they do tend to lead business cycles, whereas sectors that use proportionately less intermediate inputs tend to lag business cycles.

Lastly, we investigate how the use of imported inputs affects business-cycle moments. We find that sectors with a high share of imported inputs are also relatively more volatile. This is in line with the notion that the price of imported inputs is likely to be more volatile since part of that is tied to exchange rate fluctuations. At the same time, we find that the share of imported inputs is not predictive of business-cycle correlations.

Goods Market Pricing Distortions

The third category of industry characteristics that we examine are those capturing goods market distortions. One measure attempts to capture the competitive pressures faced by firms in different industries, the idea being that firms in more concentrated industries have more scope to vary their markups over the business cycle. The second one is a measure of nominal stickiness based on microeconomic price data. Bils et al. (2014) have defended time-varying goods market distortions as a

[4] For brevity, we do not report the numerical results for these exercises. The removed sectors are i) electric power generation, transmission and distribution; ii) oil and gas extraction; iii) natural gas distribution; and iv) petroleum and coal manufacturing.

Figure 3 Goods Market Distortion Correlates

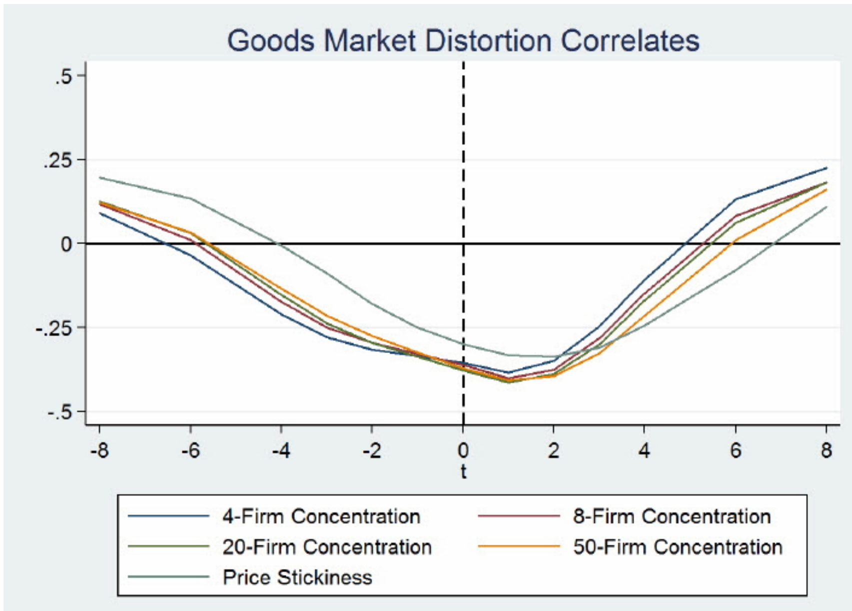

Note: Figures report the correlations between industry characteristics and business-cycle moments (either relative volatility or business-cycle correlation for various industry leads/lags).

key element in business-cycle propagation. As for nominal rigidities, they of course underlie a large literature on monetary policy and business cycles. To measure those, we use the average frequency of price adjustment as measured in the CPI data by Nakamura and Steinsson (2008).

We first examine how market concentration in different industries is related to their business-cycle behavior. We measure market concentration by the share of the top four firms in each industry. This provides a measure of the potential role for goods market pricing distortion under the assumption that firms in more concentrated industries have more scope for markup variation. We find that firms in more concentrated industries are also less cyclical.

Table 4 Goods Market Distortion Correlates

	Std. Dev.	t-8	t-4	t-3	t-2	t-1	t	t+1	t+2	t+3	t+4	t+8
Four-Firm Concentration	0.11	0.09	-0.21	-0.28	-0.32	-0.33	-0.36	-0.38	-0.35	-0.25	-0.11	0.23
Eight-Firm Concentration	0.07	0.12	-0.17	-0.25	-0.3	-0.33	-0.36	-0.4	-0.38	-0.28	-0.15	0.18
20-Firm Concentration	0.02	0.13	-0.15	-0.24	-0.29	-0.34	-0.38	-0.41	-0.39	-0.3	-0.17	0.18
50-Firm Concentration	-0.01	0.12	-0.14	-0.21	-0.27	-0.32	-0.37	-0.41	-0.39	-0.33	-0.21	0.16
Price Stickiness	-0.08	0.2	-0.01	-0.09	-0.18	-0.25	-0.30	-0.33	-0.34	-0.31	-0.24	0.11

We then examine the correlation of business-cycle statistics with the average frequency of price changes. The data indicate that industries with less sticky prices (higher frequency of price adjustment) are less correlated with business cycles. This is in line with the view that nominal rigidities play a role in the propagation of business-cycle shocks.

Financial Sensitivity

The last category we measure includes the industry characteristics that are likely to be correlated with their sensitivity to financial shocks. The most prominent one is average firm size, proposed by Gertler and Gilchrist (1994), under the idea that smaller firms are more likely to be financially constrained. We also examine firm age and a financial frictions index proposed by Hadlock and Pierce (2010) using both size and age. Two further measures of financial sensitivity are external financial dependence, proposed by Rajan and Zingales (1998) to study the role of financial development in growth, and the inventory/sales ratio, used by Schwartzman (2014), Raddatz (2006), and others to study the impact of financial shocks in less-developed economies.

We find that industries with smaller firms (and, presumably, facing higher financial frictions) tend to lag business cycles by about three quarters, but even there, the correlation is relatively moderate. On the other hand, older firms (which presumably face lower financial frictions) tend to lead business cycles. The net effect is that the size-age index implies that industries in which financial constraints are less severe lead business cycles, whereas those where they are more severe lag business cycles. This pattern does not suggest a simple story of financial frictions amplifying business cycles, but it does suggest some possibly interesting implications for the role of financial frictions in their propagation. Of course, this interpretation presumes that susceptibility to financial constraints is the major difference between firms of different ages and sizes. Presumably, those characteristics might be correlated with many other aspects of firm behavior.

A similar pattern is apparent when we use external financial dependence as a measure of sensitivity to financial conditions. External financial dependence is equal to one minus the median ratio between cash flow and capital expenditures for firms within an industry. It measures how much firms need to raise over and above their internally generated cash flow in order to finance their typical investment. We find that fluctuations in industries in which firms are more dependent on external finance are more likely to lag fluctuations in output.

Figure 4 Financial Correlates

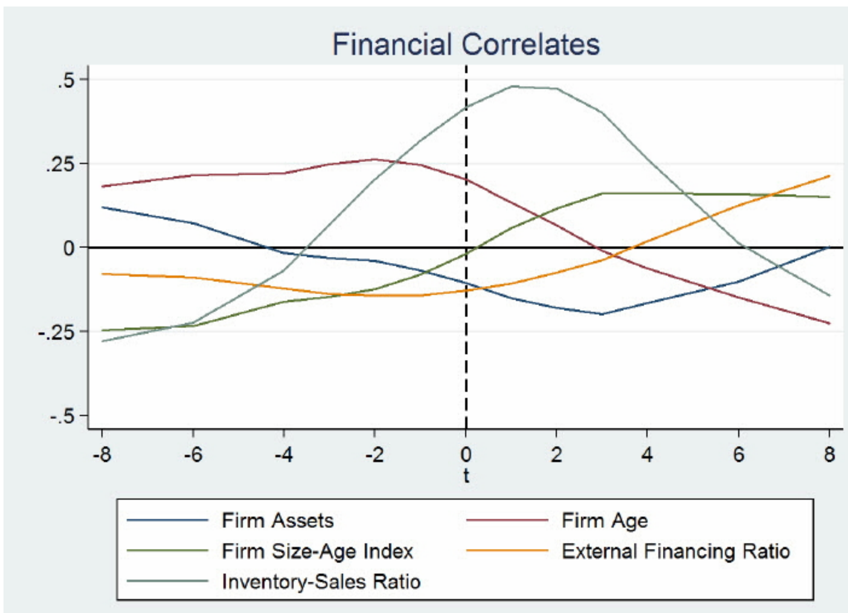

Note: Figures report the correlations between industry characteristics and business-cycle moments (either relative volatility or business-cycle correlation for various industry leads/lags).

Table 5 Financial Correlates

	Std. Dev.	t-8	t-4	t-3	t-2	t-1	t	t+1	t+2	t+3	t+4	t+8
Assets	-0.13	0.12	-0.02	-0.03	-0.04	-0.07	-0.11	-0.15	-0.18	-0.2	-0.17	0.00
Age	-0.12	0.18	0.22	0.25	0.26	0.25	0.20	0.13	0.07	-0.01	-0.06	-0.22
Size-Age Index	0.12	-0.25	-0.16	-0.15	-0.12	-0.08	-0.02	0.06	0.12	0.16	0.16	0.15
Ext. Fin. Ratio	-0.08	-0.08	-0.12	-0.14	-0.14	-0.14	-0.13	-0.11	-0.07	-0.04	0.02	0.21
Cash Flow	-0.08	0.24	-0.14	-0.2	-0.23	-0.26	-0.27	-0.29	-0.26	-0.2	-0.12	0.13
Capital Exp.	0.01	0.26	0.01	-0.06	-0.1	-0.15	-0.20	-0.26	-0.29	-0.28	-0.22	-0.03
Invt.-Sales Ratio	0.25	-0.28	-0.07	0.07	0.2	0.32	0.42	0.48	0.47	0.4	0.26	-0.14

The final industry characteristic we examine is the inventory/sales ratio. In contrast to the other measures, the business-cycle correlation for firms with a high inventory/sales ratio is fairly large. It is also particularly pronounced contemporaneously, although the peak occurs at one- or two-quarter lags.

Multivariate Analysis

The analysis so far is based off the comparison of business-cycle moments across industries taking one industry characteristic at a time. To disentangle those, we turn now to multivariate analysis, i.e., we run a simple OLS regression with the different business-cycle statistics as a dependent variable and all the industry characteristics that we explored on the right-hand side. Here we use the measure of energy intensity after excluding the four outlying sectors. This sharpens the interpretation of the results since, as pointed out by Bils et al. (2013), those very high energy intensity sectors are also sectors with very flexible prices, leading to a strong multicolinearity between energy intensity and frequency of price changes. This multicolinearity problem is eliminated once we exclude those outliers. Tables 6 and 7 present the results for the different statistics, with coefficients that are significant at a 10 percent level marked in bold. Before running the regression, all right-hand-side variables were normalized by their standard deviation, so the coefficients can be interpreted as the effect of a one standard deviation change in the value of those regressors on the various business-cycle statistics. Focusing on these statistically significant coefficients, we obtain the following results, which are robust to the introduction of multivariate controls:

1) *Volatility is higher in sectors with durable goods, imported inputs, and high frequency of price adjustment.*

The findings for durable goods and imported inputs conform to the findings from the univariate analysis above. The correlation with frequency of price adjustment only emerges in the context of the multivariate analysis. It conforms to the notion that, all else constant, firms in industries that are subject to more variable shocks will choose to adjust prices more frequently.

2) *The sectors least correlated with aggregate GDP are those producing necessities (low Engel elasticity), those that have their production oriented toward government consumption, and those that intensively use agricultural and mining inputs. Sectors oriented toward the production of intermediate inputs are more correlated with output.*

Table 6 Regression Coefficients (1)

	Std. Dev.	t-8	t-6	t-4	t-3	t-2	t-1
Four-Firm Concentration Ratio	-0.117	0.002	-0.006	-0.032	-0.043	-0.044	-0.034
Durability	**0.655**	-0.051	**-0.068**	-0.058	-0.053	-0.046	-0.033
Energy Inputs	-1.663	-0.065	-0.056	-0.079	-0.099	-0.108	-0.093
Ext. Fin. Ratio	-0.126	-0.005	-0.009	-0.024	-0.029	-0.03	-0.032
Household Share	-0.137	0.033	**0.049**	**0.048**	0.035	0.016	-0.01
Government Share	**-0.382**	-0.037	**-0.060**	**-0.083**	**-0.089**	**-0.093**	**-0.090**
Construction Share	-0.406	0.015	0.03	0.045	**0.058**	**0.073**	**0.086**
Inv.-Sales Ratio	0.077	-0.043	**-0.058**	**-0.061**	-0.047	-0.031	-0.01
Median Assets	0.397	-0.022	-0.033	-0.018	-0.001	0.016	0.021
Median Age	-0.357	0.021	0.031	0.025	0.027	0.029	0.03
Engel Curve	0.004	-0.002	0.044	0.063	0.071	**0.083**	**0.099**
Agricultural Inputs	-0.373	0.026	0.023	-0.035	**-0.062**	**-0.080**	**-0.094**
Mining Inputs	0.443	**0.125**	0.087	0.048	0.018	-0.015	-0.047
Intermediate Inputs	-0.128	0.036	0.057	**0.078**	**0.078**	**0.073**	0.062
Imported Inputs	**0.980**	0.037	0.056	0.056	0.056	0.058	0.051
Capital Share	-0.034	0.015	-0.004	-0.005	-0.007	-0.008	-0.017
Price Stickiness	**0.971**	0.037	0.038	0.057	0.074	0.076	0.07

Note: Tables report OLS coefficients for business-cycle moments against the set of industry characteristics. Coefficients significant at the 10 percent level are in bold. Each column is a separate regression.

The multivariate analysis suggests that the low correlation of sectors intensive in agricultural inputs is not a simple artifact of those sectors also being oriented toward household consumption.

The last two facts concern the dynamic relationships between sectoral output and aggregate output:

3) *Sectors that are oriented toward the private sector (have a low government share), that sell a large fraction of their output as intermediate inputs, use fewer agricultural inputs, use intermediate inputs intensively, adjust prices frequently, and are not dependent on external finance tend to lead business cycles.*

and

4) *Sectors that are government-oriented, sell a small fraction of their output as intermediate inputs, are not intensive in mining inputs, adjust prices less frequently, and are more dependent on external finance tend to lag business cycles.*

Table 7 Regression Coefficients (2)

	t	t+1	t+2	t+3	t+4	t+6	t+8
4-firm Concentration Ratio	-0.031	-0.035	-0.027	-0.01	0.012	0.046	0.051
Durability	-0.021	-0.013	-0.007	0.001	0.015	0.046	0.054
Energy Inputs	-0.067	-0.018	0.036	0.087	0.136	**0.182**	**0.150**
Ext. Fin. Ratio	-0.028	-0.026	-0.018	-0.008	0.008	0.03	0.035
Household Share	-0.034	**-0.058**	**-0.068**	**-0.061**	-0.042	-0.005	0.002
Government Share	**-0.075**	**-0.054**	-0.03	-0.001	0.029	**0.061**	**0.063**
Construction Share	**0.082**	**0.058**	0.033	0.013	0.001	-0.011	-0.01
Inv. Sales Ratio	0.018	0.037	0.045	0.043	0.036	0.02	0.011
Median Assets	0.024	0.028	0.022	0.008	-0.003	-0.027	-0.026
Median Age	0.026	0.013	0.005	0.003	0.005	0.006	-0.01
Engel Curve	**0.108**	**0.095**	0.063	0.034	0.009	-0.036	-0.06
Agricultural Inputs	**-0.086**	**-0.065**	-0.039	-0.012	0.011	0.022	0.044
Mining Inputs	-0.086	**-0.131**	**-0.162**	**-0.177**	**-0.172**	-0.117	-0.069
Intermediate Inputs	0.036	-0.002	-0.039	-0.059	-0.059	-0.049	-0.061
Imported Inputs	0.050	0.04	0.032	0.018	-0.002	-0.055	**-0.085**
Capital Share	-0.019	-0.019	-0.019	-0.015	-0.012	-0.009	-0.012
Price Stickiness	0.048	0.005	-0.048	-0.095	**-0.135**	**-0.146**	-0.082

Note: Tables report OLS coefficients for business-cycle moments against the set of industry characteristics. Coefficients significant at the 10% level are in bold. Each column is a separate regression.

Those two latter sets of facts add some interesting details to the first two. For example, it becomes clear that having demand oriented toward government consumption does not insulate a sector's output from business cycles but rather leads it to react with a lag. It is also interesting to note that sectors that are very integrated in the production chain (in the sense of using intermediate inputs intensively) tend to lead business cycles, whereas those that do not use as many intermediate inputs tend to lag. The relatively low correlation of sectors with high financial dependence also hides the fact that they respond with a lag. Finally, the regressions also point to an early response of flexible price sectors and a delayed response of sticky price sectors.

3. CONCLUSION

We asked a simple question: How do business-cycle statistics vary with sectoral characteristics? Some of the answers were predictable, others less so. The results highlight the promise and pitfalls of using industry-level data to identify driving forces and propagation mechanisms in business cycles. On the one hand, the results help focus the analysis on channels that are more likely to be relevant and take away from others that do not appear so relevant. For example, the analysis points to pricing and financial frictions as channels worth investigating but provides very little evidence of a prominent role for oil shocks. On the other hand, the results highlight the need to interpret results with care, since differences in business-cycle behavior between industries may be dominated by differences in durability or demand composition that may be correlated with other characteristics of interest.

REFERENCES

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The Network Origins of Aggregate Fluctuations." *Econometrica* 80 (September): 1977–2016.

Baxter, Marianne, and Robert G. King. 1999. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *Review of Economics and Statistics* 81 (November): 575–93.

Bernanke, Ben, and Mark Gertler. 1989. "Agency Costs, Net Worth, and Business Fluctuations." *American Economic Review* 79 (March): 14–31.

Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2013. "Testing for Keynesian Labor Demand." In *NBER Macroeconomics Annual*, vol. 27, eds. Daron Acemoglu, Jonathan Parker, and Michael Woodford. Chicago: University of Chicago Press, 311–49.

Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2014. "Resurrecting the Role of the Product Market Wedge in Recessions." Working Paper 20555. Cambridge, Mass.: National Bureau of Economic Research. (October).

Gertler, Mark, and Simon Gilchrist. 1994. "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms." *Quarterly Journal of Economics* 109 (May): 309–40.

Hadlock, Charles J., and Joshua R. Pierce. 2010. "New Evidence on Measuring Financial Constraints: Moving Beyond the KZ Index." *Review of Financial Studies* 23 (May): 1909–40.

Hamilton, James D. 1994. *Time Series Analysis.* Princeton: Princeton University Press.

Hamilton, James D. 2003. "What is an Oil Shock?" *Journal of Econometrics* 113 (April): 363–98.

Hodrick, Robert J., and Edward C. Prescott. 1997. "Postwar U.S. Business Cycles: An Empirical Investigation." *Journal of Money, Credit, and Banking* 29 (February): 1–16.

Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit Cycles." *Journal of Political Economy* 105 (April): 211–48.

Nakamura, Emi, and Jón Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123 (November): 1415–64.

Prescott, Edward C. 1986. “Theory Ahead of Business-Cycle Measurement.” *Carnegie-Rochester Conference Series on Public Policy* 25 (Autumn): 11–44

Raddatz, Claudio. 2006. “Liquidity Needs and Vulnerability to Financial Underdevelopment.” *Journal of Financial Economics* 80 (June): 677–722.

Rajan, Raghuram G., and Luigi Zingales. 1998. “Financial Dependence and Growth.” *American Economic Review* 88 (June): 559–86.

Ramey, Valerie A. 2011. “Can Government Purchases Stimulate the Economy?” *Journal of Economic Literature* 49 (September): 673–85.

Rotemberg, Julio J., and Michael Woodford. 1999. “The Cyclical Behavior of Prices and Costs.” In *Handbook of Macroeconomics,* vol. 1, part B, eds. John B. Taylor and Michael Woodford. Amsterdam: North-Holland: 1051–135.

Schwartzman, Felipe. 2014. “Time to Produce and Emerging Market Crises.” *Journal of Monetary Economics* 68 (November): 37–52.

APPENDIX

A.1 List of Industries

Our industrial classification is primarily based on the four-digit 2007 NAICS codes, with certain four-digit industries consolidated into a single category to facilitate the construction of either the PCE/industry crosswalk or the industry controls. The full list of industries used is displayed in Table 8.

A.2 PCE/Industry Crosswalk

We use the 2007 PCE Bridge Table published by the Bureau of Economic Analysis to match PCE expenditure categories to industries. The Bridge Table contains consumer spending levels by PCE category and commodity pairs. For each pair, the level of total spending going to producers, wholesalers, retailers, and transport is provided.

Also included in the Bridge Table file is a concordance of commodity categories with NAICS codes. This allows the commodities to be matched with our industry groups. However, this concordance is less granular in many cases than our industry classification—these industries are pooled for the purposes of constructing the PCE/industry crosswalk. In addition, while the analysis in the paper focuses predominantly on manufacturing and related sectors, for the purpose of constructing this crosswalk, it is important to capture all sectors of the economy in order to construct a more detailed PCE/industry crosswalk. For this purpose, we make use of all the commodities and industries present in the Bridge Table.

Using this commodity/NAICS concordance with the expenditure data in the bridge table, we obtain expenditure estimates for producer margins by PCE category and industry. No observations for the wholesale, retail, or transport industries exist, as their expenditure is contained in the corresponding wholesale, retail, and transport margins for each PCE/industry pair. To create observations for these industries, we total the entire margin across a given PCE category and use this total as the value for that PCE/industry category pair. For instance, we total all wholesale margins across the "auto leasing" PCE category, and this is taken as the value for the wholesale/auto leasing pair. We do this for each PCE category and for wholesale, retail, and transport.

For these, we sum the total wholesale margin across all industries for a given PCE category and construct an additional observation designating the total as the expenditure for a given PCE category and wholesale industry pair. We repeat this process for all PCE categories and do the same for retail and transport as well.

Given total consumer expenditure broken down by PCE category/industry pairs, we construct a crosswalk between the two categories using expenditure share weights. This allows the translation of some set of values at the PCE level to the industry level, or vice versa. For each PCE/industry pair, the crosswalk contains two weights; one is the proportion of the total industry expenditure that is also from the PCE category, and the other is the proportion of total PCE category expenditure that is also from the industry. The former is used to translate PCE-level data to the industry level and the latter from the industry to the PCE level.

As an example of how this occurs, consider a set of data at the industry level with one value per industry. This dataset is merged with the crosswalk so that now each PCE category/industry pair contains both the expenditure-share weights and the industry-level data value. The PCE-level data are then estimated as the weighted average for the PCE category across all industries. This provides an estimate of the PCE value by imputing the data from the constituent industries that make up the PCE category. Using the other weight that exists for each PCE category/industry pair, the same process can occur in reverse, with PCE data translated to the industry level.

Note that, as stated above, some industries do not have a unique commodity code in the original Bridge Table and were thus pooled for the construction of the crosswalk. For these industry groups, the crosswalk will provide a single value for the group rather than a separate value for each industry. In these cases, we assume that all industries share this value in common.

A.3 Controls

A.3.1 Concentration Ratios

Industry-concentration data are taken from the 2007 Economic Census. For each 2007 NAICS industry at the six-digit level, the census contains the percentage of total industry sales from the largest four, eight, twenty, and fifty firms, along with total industry revenue. We match each six-digit NAICS category to the industry in which it is contained and take the revenue-weighted mean across all six-digit NAICS within the industry as the concentration ratio for that industry. This provides a four-firm, eight-firm, twenty-firm, and fifty-firm concentration ratio for each of our industries.

For robustness, we construct additional concentration measures from the same data: in addition to taking the revenue-weighted mean, we also take both the median and the maximum concentration ratio across six-digit NAICS industries. This leaves us with twelve values, corresponding to either four, eight, twenty, or fifty-firm concentration ratios, and to either the mean, median, or maximum across subindustries.

A.3.2 Durability

The BEA publishes depreciation/durability estimates for consumer durables, equipment, and structures. We match each PCE category to a durable good, equipment, or structure category if a corresponding category exists. We then take the service life estimate published by the BEA as a measure of the durability of the item. Nondurable goods are assigned a durability of zero. Values are then translated to the industry level using the PCE category/industry crosswalk.

A.3.3 Inputs

From the 2007 Benchmark Input/Output Use Table, we calculate the exposure of an industry to energy, agriculture, mining, as well as the industry's use of intermediate inputs. Using the commodity/NAICS crosswalk provided with the Use Table, we match each commodity to its corresponding industry and aggregate the Use Table to our industry classification. Where the provided concordance is not granular enough for our industry classification, we pool industries and assign the corresponding values to all industries in the group.

Energy Inputs: We take the proportion of total intermediate inputs that are from (1) electrical power generation, (2) oil and gas extraction, (3) natural gas distribution, and (4) petroleum and coal products manufacturing as a measure of each industry's energy exposure.

Agricultural Inputs: We take the proportion of total intermediate inputs that are from (1) crop production, (2) animal production and aquaculture, and (3) support activities for agriculture and forestry as a measure of each industry's exposure to agriculture.

Mining Inputs: We take the proportion of total intermediate inputs that are from (1) metal ore mining and (2) nonmetallic mineral mining and quarrying as a measure of each industry's exposure to mining.

Total Intermediate Inputs: We construct a measure of the total intermediate inputs used by the industry by taking the ratio of all industry inputs to the industry's output.

A.3.4 Capital Share

Also from the Use Table, we estimate the relative intensity of capital as opposed to labor in each sector. As for the input measures, we first aggregate the Use Table to our industrial classification. To do so, we compute the ratio of gross operating surplus over the sum of gross operating surplus and compensation to employees.

A.3.5 Output Shares

Again from the Use Table, we estimate several measures related to the destination of each industry's output. As before, we aggregate the Use Table to our industry classification.

Household Output: We calculate the household share as the proportion of industry output that goes to PCE.

Government Output: We calculate the government share as the total output sold to all federal, state, and local government categories listed in the Use Table as a ratio to total industry output.

Construction Output: We calculate the construction share as the proportion of each industry's output that is purchased by the construction sector.

Total Intermediate Output: We construct the proportion of total industry output that was used as an intermediate inputs by any other industry. For robustness, we also take the raw number of intermediate inputs sold without normalizing by industry output.

A.3.6 Imports and Exports

The Use Table also contains information on imports and exports by industry and can therefore also be used to calculate several measures describing the international linkages of each sector.

Import Penetration: For each industry, we take the value of industry outputs that are imported into the United States and divide by total industry production plus imports minus exports. This provides the share of each industry's final goods sold domestically that were produced internationally.

Exports: We calculate the export ratio as the share of industry output that is exported.

Imported Inputs: To measure the level of input connections to foreign markets, we calculate the ratio between imported intermediate inputs to total industry output.

Imported Share of Inputs: As an alternative measure of the input connections to foreign markets, we calculate the ratio of the total industry inputs that are imported.

A.3.7 External Financing Ratio, Cash Flow, and Capital Expenditure

Using capital expenditure and cash flow by firm and year from Compustat for 1979 through 2015, we can construct the external financing ratio as in Rajan and Zingales (1998), as one minus the ratio between cash flow to capital expenditure. Matching each firm to an industry, we take the median capital expenditure value across firms for each industry and year. Then, we take the median again across years to obtain a single value for each industry. The same procedure is used to obtain a median cash flow and median capital expenditure value for each industry. Rajan and Zingales (1998) describe the construction of the cash flow variable in greater detail.

A.3.8 Inventory Sales Ratio

From Compustat we take firm-level data on annual inventories and total sales from 1979 through 2015. From this, we normalize inventories by total sales for each firm. Matching firms to industries, we then take the median value for each industry and year and then select the median across years as the final industry value.

A.3.9 Size-Age Index

To construct measures of industry-specific financial constraints, we follow Hadlock and Pierce (2010), who show that an index that is linear in firm age and quadratic in firm asset size can capture the degree of firm financing constraints. Specifically, the index is calculated as $-.737 * size + .043 * size^2 - .04 * age$. We calculate this index for each firm and year between 1979 and 2015. Matching firms to industries, we take the median for each industry/year pair and again for each industry. We do the same for asset size and age separately.

A.3.10 Luxury Goods

We construct two measures of the degree to which the outputs of each industry are luxury goods. First, we use BLS data from the Consumer Expenditure Survey, which details the consumption expenditures for various goods by income decile. Matching these expenditure categories with PCE categories, we construct estimates of expenditures for each PCE category for the fourth and sixth income deciles and take the ratio of these values as an estimate of the luxury status of a PCE category. We then use the PCE/industry crosswalk to map these values to the industry level.

As an alternate measure of the income elasticity of industry output, we also take the Engel Curve slopes estimated by Bils et al. (2013). They estimate these Engel Curve values for PCE categories, which we map into the industry level using our PCE/industry crosswalk.

A.3.11 Price Stickiness

To capture the frequency of price changes within an industry, we take the price-adjustment durations estimated by Nakamura and Steinsson (2008). The estimates are provided at the Entry Line Item (ELI) level. By using the ELI/PCE crosswalk provided by the BLS, we can transfer these ELI-level duration values to the PCE classification. For each PCE category, we assign the average of the duration values for the set of ELIs with which the PCE category is matched. Following this, we can match PCE-level values to the industry level using the PCE/industry crosswalk.

Table 8 Industries

Industry	2007 NAICS
Oil and gas extraction	211-
Coal mining	2121
Metal ore mining	2122
Nonmetallic mineral mining and quarrying	2123
Support activities for mining	213-
Electric power generation, transmission, distribution	2211
Natural gas distribution	2212
Animal food manufacturing	3111
Grain and oilseed milling	3112
Fruit and vegetable preserving and specialty food manufacturing	3114
Dairy product manufacturing	3115
Animal slaughtering and processing	3116
Bakeries and tortilla manufacturing	3118
Other food manufacturing	3119
Beverage manufacturing	3121
Tobacco manufacturing	3122
Textile mills and textile product mills	313-, 314-
Apparel, leather, and allied manufacturing	315-, 316-
Sawmills and wood preservation	3211
Veneer, plywood, engineered wood product manufacturing	3212
Other wood product manufacturing	3219
Pulp, paper, and paperboard mills	3221
Converted paper product manufacturing	3222
Printing and related support activities	323-
Petroleum and coal products manufacturing	324-
Basic chemical manufacturing	3251
Resin, synthetic rubber, artificial synthetic fibers and filaments manufacturing	3252
Pesticide, fertilizer, other agricultural chemical manufacturing	3253
Pharmaceutical and medicine manufacturing	3254
Paint, coating, and adhesive manufacturing	3255
Soap, cleaning compound, and toilet paper manufacturing	3256
Plastics product manufacturing	3261
Rubber product manufacturing	3262
Clay product and refractory manufacturing	3271
Glass and glass product manufacturing	3272
Cement and concrete product manufacturing	3273
Lime, gypsum and other nonmetallic mineral product manufacturing	3274, 3279
Alumina and aluminum production and processing	3313
Nonferrous metal (except aluminum) production and processing	3314
Foundries	3315
Forging and stamping	3321
Cutlery and handtool manufacturing	3322
Architectural, construction, and mining machinery manufacturing	3323
Hardware manufacturing	3325
Spring and wire product manufacturing	3326
Machine shops, turned product, screw, nut, bolt manufacturing	3327
Coating, engraving, heat treating, and allied activities	3328
Other fabricated metal product manufacturing	3329
Agricultural, construction, and mining machinery manufacturing	3331
Industrial machinery manufacturing	3332

Table 8 (Continued) Industries

Ventilation, heating, air conditioning, and commerical refrigeration equipment manufacturing	3334
Metalworking machinery manufacturing	3335
Engine, turbine, power transmission equipment manufacturing	3336
Computer and peripheral equipment manufacturing	3341
Communications equipment manufacturing	3342
Audio and video equipment manufacturing	3343
Semiconductor & other electronic component manufacturing	3344
Navigational, measuring, electromedical, and control instruments manufacturing	3345
Electric lighting equipment manufacturing	3351
Household appliance manufacturing	3352
Electrical equipment manufacturing	3353
Other electrical equipment and component manufacturing	3359
Motor vehicle manufacturing	3361
Motor vehicle body and trailer manufacturing	3362
Motor vehicle parts manufacturing	3363
Aerospace product and parts manufacturing	3364
Railroad rolling stock manufacturing	3365
Ship and boat building	3366
Other transportation equipment manufacturing	3369
Household and institutional furniture and kitchen cabinet manufacturing	3371
Medical equipment and supplies manufacturing	3391
Newspaper, periodical, book, and directory publishers	5111