



DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

EMPLOYMENT CONCENTRATION ACROSS US COUNTIES

Klaus Desmet and Marcel Fafchamps

Number 180

December 2003

Manor Road Building, Oxford OX1 3UQ

Employment Concentration across US Counties

Klaus Desmet
Universidad Carlos III de Madrid
and CEPR*

Marcel Fafchamps
University of Oxford[†]

October 2003
First draft

Abstract

This paper examines the spatial distribution of jobs across US counties and investigates whether sectoral employment is becoming more or less concentrated. The existing literature has found deconcentration (convergence) of employment across urban areas. Cities only cover a small part of the US, though. Using county data, our results indicate that deconcentration is limited to the upper tail of the distribution. The overall picture is one of increasing concentration (divergence). While this seemingly contradicts the well documented deconcentration in manufacturing, we show that these aggregate employment dynamics are driven by services. Non-service sectors – such as manufacturing and farming – are indeed becoming more equally spread across space, but services are becoming increasingly concentrated.

*Department of Economics, Universidad Carlos III de Madrid, 28903 Getafe (Madrid), Spain. E-mail: desmet@eco.uc3m.es.

[†]CSAE, Department of Economics, University of Oxford, Manor Road, Oxford OX1 3UQ, United Kingdom. E-mail: marcel.fafchamps@economics.oxford.ac.uk.

1 Introduction

Economic activity is unevenly distributed across space. The interaction of positive and negative externalities create intricate geographical patterns of city clusters and rural hinterland (Henderson, 1988; Fujita, Krugman and Venables, 1999). Over time, these patterns evolve because of changes in preferences, production technologies and transport costs. As a result, the spatial distribution of employment evolves as jobs are created in certain locations, and destroyed elsewhere. Understanding how economic activity is likely to be distributed through space in the future is of utmost importance for policy makers at the national and local level.

This paper describes the geographical evolution of jobs in the US between 1972 and 1992, with the goal of understanding what the future spatial distribution of employment will look like if current tendencies continue. We use county-level employment in 13 different sectors — ranging from farming to manufacturing and services — and focus on the ergodic distribution of jobs.

Our work differs from the existing literature in many respects. First, rather than looking at income per capita or population, we are interested in employment. Many authors have studied whether standards of living in the US are becoming more similar over time. For instance, Higgins, Levy and Young (2003) find strong evidence of income convergence across counties. Given the high degree of labor mobility (Blanchard and Katz, 2003), this is not entirely surprising. However, income convergence does not tell us anything about where economic activity is locating. Is the US moving towards a situation with more or with less large and medium-sized cities? Are rural counties losing or gaining jobs? These are the kind of questions we can address in our paper.¹ This seems similar to studying the distribution of population, a topic that has drawn a lot of attention, especially from the angle of trying to document Zipf’s law for cities (see, e.g., Eaton

¹When comparing countries with between which labor and capital are not fully mobile, the distribution of GDP per capita can be regarded as capturing the distribution of economic activity across space. However, in a country like the US, where capital and workers are highly mobile, the dispersion of GDP per worker across geographical units is more a measure of dispersion in productivity than in economic activity per se.

and Eckstein, 1997). However, looking at employment, rather than population, has the advantage of allowing sectoral disaggregation.

Second, we examine the country as a whole, not just metropolitan areas. Most of the literature on the spatial organization of economic activity in the US has focused on cities. This is probably because some of the more striking trends – such as suburbanization – have been affecting metropolitan areas. As a result, counties are somewhat understudied. However, as pointed out by Beeson, DeJong and Troesken (2001), limiting the analysis to urban areas introduces a selection bias, since cities are those areas which experienced high growth in the past. By studying counties, we do not have this problem, as we are looking at the entire distribution, and not just at its upper tail.

Our third, and perhaps most important point of departure with the existing literature, is our methodology. Instead of relying on a single method — whether β convergence, σ convergence, or ergodic distributions — we develop a methodology that combines them all, in which all of them are nested. Much of the existing work comparing geographical units is couched in terms of Barro's β convergence: the underlying model is deterministic in nature (Barro, 1991; Mankiw, Romer and Weil, 1992). As first emphasized by Quah, evidence of β convergence can yield a misleading picture, as it can arise even when countries or regions are getting further apart, and vice versa (Quah, 1993; Quah, 1996a; Durlauf and Quah, 1999). As a solution, Sala-i-Martin (1996) suggests studying distributions by looking at the evolution of the variance over time, a concept known as σ convergence. Quah (1996b, 1997) proposes instead to focus on the ergodic distribution. This refers to the long term spatial distribution of economic activity that would arise if current transition probabilities would remain constant.

In this paper we compute parametric and non-parametric versions of the unconditional and conditional β -convergence tests and explain how they can be understood as describing the expectation of the transition probability. We also compute two versions of the ergodic distribution. The first version is the stochastic equivalent of unconditional β convergence: it assumes that all counties are inherently equivalent and could switch place with each other over time. The second version is the stochastic equivalent of conditional

β convergence: it conditions on county characteristics that are constant over time. It is our best estimate of how economic activity would be distributed across US counties should current tendencies remain unchanged.

We also introduce a number of practical innovations when deriving the ergodic distribution. These innovations make the results both more detailed and more accurate.² They are easy to implement and are applicable to any empirical investigation involving distribution dynamics.

We now turn to describing our main findings. Whereas recent work on metropolitan areas shows a tendency towards deconcentration, with total employment becoming more equally spread across *cities* (Chatterjee and Carlino, 2001), standard β convergence tests using US *county* data suggest the contrary, with jobs becoming more concentrated over recent decades (Desmet and Fafchamps, 2003). The analysis presented here resolves this apparent puzzle. Results show that, compared to the current distribution of total employment across counties, the ergodic distribution is a lot flatter, with the middle group thinning out. The overall picture that emerges is thus one of concentration (divergence), with lots of small and medium sized counties losing jobs to the more urban ones. At the upper tail of the distribution, however, the opposite is true, with large cities losing jobs in favor of intermediate sized urban areas. In other words, there is deconcentration (convergence) in the upper part of the distribution, and concentration (divergence) in the distribution at large. This explains the opposing results of Carlino and Chatterjee (2001) and Desmet and Fafchamps (2003).

The increased concentration evident in total employment stands in contrast to what happened within the manufacturing sector. It is by now a stylized fact that since World War II manufacturing employment has become less concentrated, albeit at a slow pace (Dumais, Ellison and Glaeser, 2002; Kim, 1995).³ Our data confirm this empirical regularity. Whereas manufacturing cannot explain the spatial dynamics of aggregate

²Moreover, contrary to Quah who uses a highly complex programming language in Unix to obtain non-parametric kernel estimates of the transition matrix, we rely on simple Stata commands. The ado files are available from the authors upon request.

³Note that this is not the same as the other well documented phenomenon of suburbanization.

employment, services can. Two of the three main service sectors — ‘retail’ and ‘finance, insurance and real estate’ — indeed exhibit concentration (divergence) in the middle part of the distribution and deconcentration (convergence) in the upper tail. This is most patent in the case of ‘finance, insurance and real estate’, where we get ‘twin peaks’ — a bimodal ergodic distribution.

That overall trends in the economy are driven by services should not come as a surprise, given their numerical importance. However, the fact that services behave differently from the rest of the economy is interesting, because much of the research in economic geography has focused on manufacturing. Our research confirms that the much heralded demise of cities, epitomized by suburbanization and weakening agglomeration economies, is not occurring. The reason is the rise of the service industry, particularly in retail trade and in finance, insurance, and real estate (Kolko, 1999).

2 Methodology

In this section we present a general framework to discuss the evolution of the distribution of an arbitrary geographical variable over time. Let our variable of interest be denoted y_t^i where i stands for location and t for time. Variable y_{it} could denote GDP per head, employment, or income, but for now this is of no importance. At each point in time we observe realizations of y_t^i for each i . We want to know whether, over time, realizations of y_t^i are becoming more ‘alike’ across all i ’s. This we call convergence. If realizations are becoming less alike, we call it divergence. We first discuss unconditional convergence; we then look at conditional convergence.

2.1 Unconditional convergence in a deterministic model

One approach to this question is that taken by the growth convergence literature (Durlauf and Quah, 1999). This literature reverts around a β -convergence test meant to ascertain whether GDP per head across countries is converging towards a common value y^* . This test is implemented via a regression of the form (Mankiw, Romer and Weil, 1992; Quah,

1993):

$$\log y_{t+1}^i - \log y_t^i = \alpha - \beta \log y_t^i \quad (1)$$

to which an error term is added for estimation purposes. Equation (1) can be rewritten as:

$$y_{t+1}^i = e^\alpha \left(y_t^i \right)^{1-\beta} \quad (2)$$

Equation (2) is a deterministic difference equation with two steady states:

$$\begin{aligned} y^* &= e^{\frac{\alpha}{\beta}} \\ y_0 &= 0 \end{aligned}$$

which, in the unconditional case, are the same for all i 's. The stability of this deterministic system around the y^* steady state depends on the sign of β : if $\beta < 0$, the y^* steady state is not stable and y_t diverges from it. Standard convergence tests estimate equation (1) and examine whether β is positive or not.

One critique of this model is that it imposes too much structure on the law of motion of y_{t+1}^i . In particular, it is unable to test for the presence of multiple steady states. It also assumes that convergence is exponential: the linear approximation underlying equation (1) is only valid locally. It makes little sense to apply this approximation to observations which, according to the model, are very far from y^* . A more satisfactory model is one that allows for non-linearity:

$$\log y_{t+1}^i = \phi(\log y_t^i) \quad (3)$$

where $\phi(\cdot)$ is an arbitrary smooth function. Equation (3) can be estimated by standard non-parametric techniques. If function $\phi(\cdot)$ cuts the 45 degree line more than once, the process driving y_t^i is shown to have multiple equilibria. Each point at which $\phi(\cdot)$ cuts the 45 degree line from above is stable; each point where it cuts from below is unstable. As it turns out, it is easier to graph the equivalent alternative model:

$$\begin{aligned} \log y_{t+1}^i - \log y_t^i &= \phi(\log y_t^i) - \log y_t^i \\ &= f(\log y_t^i) \end{aligned} \quad (4)$$

Estimates of equation (4) are presented in the empirical section. Evidence of multiple deterministic steady states is found for several sectors.

2.2 Unconditional convergence in a stochastic model

As Durlauf and Quah (1999) have emphasized, the approach to convergence based on a deterministic model is hardly appropriate because it fails to recognize that in practice y_{it} is stochastic. A more satisfying approach is to regard y_{it} as stochastic. To illustrate this point, let us return for a moment to the linear model, to which we add a stochastic term u_t^i :

$$\log y_{t+1}^i - \log y_{it} = \alpha - \beta \log y_t^i + u_t^i \quad (5)$$

In this case, y_{it} never actually settles anywhere permanently so there is no steady state in the deterministic sense and thus no β -convergence. As Quah (1993) has shown, β in this context measures the extent to which realizations of y_{it} are correlated over time.

When y_{it} is stochastic, a more adequate representation of its evolution over time is:

$$f_{t+1}(y_{t+1}) = \int_{-\infty}^{\infty} g(y_{t+1}|y_t) f_t(y_t) \quad (6)$$

where $f_t(y_t)$ denotes the (unconditional) distribution of y_t at time t across all i 's and $g(y_{t+1}|y_t)$ denotes its transition probability. Here and in the remainder of this section, we assume the transition probability to be constant over time. This is of course an oversimplification. We revisit this issue in the empirical section. Equation (6) is itself a deterministic law of motion. Provided that certain conditions are satisfied (Stokey, Lucas, Prescott, 1989; Luenberger, 1979), this system has a steady state or time-invariant distribution $f(y_{t+1})$ to which it converges.⁴ This time-invariant distribution is called the ergodic distribution. It is the distribution $f(y_{t+1})$ that satisfies:⁵

$$f(y_{t+1}) = \int_{-\infty}^{\infty} g(y_{t+1}|y_t) f(y_t) \quad (7)$$

⁴In empirical applications, the most important issue that arises with respect to the existence of a non-degenerate ergodic distribution is that of de-trending. This is discussed in detail below.

⁵There might be multiple solutions to equation (ergodic) and thus multiple ergodic distributions. Multiple solutions do not arise in our empirical analysis and are ignored here.

β -convergence corresponds to the case when the ergodic distribution is degenerate with a mass point at y^* . In general, $f(y_{t+1})$ is not degenerate. If the ergodic distribution $f(y_{t+1})$ is more concentrated — has lower variance — than the current distribution $f_t(y_t)$, we conclude that there is convergence, and vice versa. This is but a straightforward extension of the concept of σ -convergence introduced by Sala-i-Martin (1996). The advantage of dealing with the ergodic distribution is that we can extrapolate current transition probabilities to the indefinite future, hence obtaining a clearer picture of what these probabilities imply for the long-term.

2.3 Empirical implementation

In empirical implementations, we begin by estimating linear and non-linear β -convergence models and look for possible evidence of multiple deterministic steady states. We then turn to the stochastic approach and derive ergodic distributions. Computing the ergodic distribution involves three steps: (1) calculating $f_{t+1}(y_{t+1})$, $f_t(y_t)$, and $f_{t+1,t}(y_{t+1}, y_t)$; (2) deriving $g(y_{t+1}|y_t)$ from the fact that the conditional distribution is the joint distribution divided by the marginal distribution:

$$g(y_{t+1}|y_t) = \frac{f_{t,t+1}(y_{t+1}, y_t)}{f_t(y_t)}$$

and (3) obtaining the ergodic distribution by solving (7). As illustrated by Quah (1996b), steps (1) and (2) are easily handled by non-parametric techniques: $f_t(y_t)$, and $f_{t+1,t}(y_{t+1}, y_t)$ are estimated by fitting a kernel density to the data, and $g(y_{t+1}|y_t)$ is obtained by dividing one by the other. For the third step, it is difficult to work with equation (7) directly. The standard approach in practice is to discretize the state space. Formally, let:

$$\begin{aligned} p_k &\equiv \Pr(y = Y_k) \\ a_{km} &= \Pr(y_{t+1} = Y_k | y_t = Y_m) \end{aligned}$$

with $k, m = \{1, \dots, N\}$ with N possible states (intervals) of y . The ergodic distribution then is the solution to:⁶

$$p_k = \sum_m a_{km} p_m$$

Because step (3) involves discretization, it is common to discretize steps (1) and (2) as well and to compute the transition matrix $[a_{km}]$ directly from the data. This approach, however, fails to take advantage of the smoothing properties of kernel densities. For this reason, transition matrices used in practice are usually very coarse – e.g., with four or five intervals only. The resulting ergodic distribution is too rough to draw precise conclusions about convergence. In this paper, we obtain a more refined estimate of the ergodic distribution by postponing discretization until step (3) and computing the transition matrix from the smoothed conditional distribution $g(y_{t+1}|y_t)$ rather than directly from the data.⁷

An important practical detail is a proper normalization of the data. Suppose y_t^i has a trend component so that it is growing over time. If $f_t(y_t)$, and $f_{t+1,t}(y_{t+1}, y_t)$ are

⁶In matrix form we have:

$$\begin{aligned} p &= Ap \\ (I - A)p &= 0 \end{aligned}$$

It looks like the above system only has a solution of the form $p = 0$, but this is an illusion. Matrix A does not have full rank since, by definition of a probability, each column sums to 1. To find p , one needs to drop one row of A and to add the requirement that:

$$\sum_i p_i = 1$$

We obtain a system of the form:

$$\begin{aligned} \begin{bmatrix} 1 - a_{11} & \dots & -a_{1N} \\ \dots & 1 - a_{ii} & \dots \\ -a_{N-1,1} & \dots & -a_{N-1,N} \\ 1 & \dots & 1 \end{bmatrix} [p] &= \begin{bmatrix} 0 \\ \dots \\ 0 \\ 1 \end{bmatrix} \\ Bp &= b \end{aligned}$$

The modified system can be solved by inverting:

$$p = B^{-1}b$$

In a linear system such as this, the ergodic distribution is in general unique cit [Luenberger dynamic].

⁷Quah (1996a, 1996b, 1997) computes the transition matrix (steps 1 and 2) using non-parametric techniques. But when computing the ergodic distribution, he uses a crudely discretized transition matrix with a small number of cells only.

estimated directly from the data, the resulting ergodic distribution will be a mass point at the highest value for y_t . This is normal: if y_t^i grows without bounds, all realizations of y_t^i will eventually exceed the highest value of y_t^i at time t . The data must therefore be detrended first before undertaking steps (1) to (3). The ergodic distribution then represents how the distribution of y_t^i around its mean will evolve over time.

The shape of the ergodic distribution is potentially affected by whether the trend is linear or exponential. With a linear trend, the data generating process can be written:

$$\begin{aligned} y_t^i &= z_t^i \\ y_{t+s}^i &= \beta s + z_{t+s}^i \end{aligned} \tag{8}$$

where βs is a common time trend. This means that irrespective of their initial value, all y_t^i 's grow by the same quantity.⁸ Taking expectations we have:

$$\begin{aligned} E[y_t^i] &= E[z_t^i] \\ E[y_{t+s}^i] &= \beta s + E[z_{t+s}^i] \end{aligned}$$

Normalization means setting $E[z_t^i] = E[z_{t+s}^i]$. This yields:

$$\beta = \frac{E[y_{t+s}^i] - E[y_t^i]}{s}$$

or a sample equivalent. The corrected data for the second period is thus:

$$\begin{aligned} z_{t+s}^i &= y_{t+s}^i - \beta s \\ &= y_{t+s}^i - (E[y_{t+s}^i] - E[y_t^i]) \end{aligned}$$

Things are slightly more complex when comparing years with different average growth rates, as when using the PWT data. If the growth parameter β is the same for all sub-periods, it makes sense to estimate it using all time periods, i.e., to use a common β for all time periods. If, however, we believe the growth rate in different sub-periods was different, then it would be more appropriate to use a different β for each period.

⁸ If $y_t^i = \log(x_t^i)$, this means that all x_t^i 's grow at the same rate.

With an exponential trend, the data generating process becomes:

$$\begin{aligned} y_t^i &= z_t^i \\ y_{t+s}^i &= z_{t+s}^i e^{\beta s} \end{aligned} \tag{9}$$

Again we want $E[z_t^i] = E[z_{t+s}^i]$. Straightforward algebra yields:

$$\begin{aligned} e^{\beta s} &= \frac{E[y_{t+s}^i]}{E[y_t^i]} \\ z_{t+s}^i &= y_{t+s}^i \frac{E[y_t^i]}{E[y_{t+s}^i]} \end{aligned}$$

Which approach is valid depends on the underlying data generation process. It is easy to characterize the bias introduced by choosing the wrong model. Suppose the true model is (8). Further suppose that there is positive growth, i.e., $(E[y_{t+s}^i] - E[y_t^i]) > 0$. Correcting using (9) yields the misspecified variable m_{t+s}^i :

$$\begin{aligned} m_{t+s}^i &= y_{t+s}^i \frac{E[y_t^i]}{E[y_{t+s}^i]} \\ m_{t+s}^i &> z_{t+s}^i \text{ if } y_{t+s}^i \text{ is small} \\ m_{t+s}^i &< z_{t+s}^i \text{ if } y_{t+s}^i \text{ is large} \end{aligned}$$

Whenever y_{t+s}^i and y_t^i are strongly correlated and y_{t+s}^i is close to y_t^i , using (9) incorrectly results in an overestimate of convergence probabilities, i.e., the transition up for low initial states and the transition down for high initial states. The opposite is also true: incorrectly applying (8) when (9) is true increases the variance of the ergodic distribution and thus underestimates convergence.

2.4 Conditional convergence

Until now we have implicitly assumed that all y_t^i have the same distribution, irrespective of their inherent, time-invariant characteristics X^i . In many situations, this is an unrealistic assumption: part of the variation in y_t^i across i 's is due to differences in their X^i . This variation does not disappear over time. Consequently, if we fail to control for X^i we may falsely conclude that y_t^i is not converging when in fact it is. Correct inference about

convergence therefore requires that we condition on X^i . We call this approach conditional convergence.

We decompose the variation in y_t^i into two parts: that due to $X^i\theta$ and that due to a first order stochastic process z_t^i . The data generation process takes the form:

$$y_t^i = X^i\theta + z_t^i$$

with $E[Xz] = 0$ and:

$$f_{t+1}(z_{t+1}) = \int_{-\infty}^{\infty} g(z_{t+1}|z_t) f_t(z_t) \quad (10)$$

with corresponding ergodic equation

$$f(z_{t+1}) = \int_{-\infty}^{\infty} g(z_{t+1}|z_t) f(z_t) \quad (11)$$

Although we do not observe z_t^i directly, we can obtain a consistent estimate by first estimating β from a pooled regression of the form

$$y_t^i = X^i\theta + z_t^i \quad (12)$$

...

$$y_{t+s}^i = X^i\theta + z_{t+s}^i$$

and then using $\hat{\beta}$ to compute

$$\hat{z}_t^i = y_t^i - X^i\hat{\theta} \quad (13)$$

...

$$\hat{z}_{t+s}^i = y_{t+s}^i - X^i\hat{\theta}$$

The ergodic distribution of \hat{z}_t^i can then be estimated through steps (1)-(3) as detailed before.

The long-term distribution of y_t^i is obtained by combining the variation due to $X^i\beta$ to the ergodic distribution of z_t^i . Let this distribution be written $f_Y(y)$. We have:

$$\begin{aligned} y^i &= X^i\hat{\theta} + z^i \\ &= w^i + z^i \end{aligned}$$

This is a standard problem in statistics. The general formula in the discrete case is (Mood, Graybill and Boes, 1974, p. 186):

$$\begin{aligned} f_Y(y) &= \sum_w f_Z(y-w)f_W(w) \\ &= \sum_{X\hat{\theta}} f_Z(y-X\hat{\theta})f_W(X\hat{\theta}) \end{aligned} \quad (14)$$

where $f_Z(\cdot)$ is the ergodic distribution function of z^i . Applying this formula to the data yields the conditional ergodic distribution of y_t^i .⁹

2.5 Ergodic distribution and β -convergence

It is useful to illustrate how our approach to conditional convergence relates to the standard β -convergence literature and how conditional convergence can be implemented in the non-linear β -convergence model. Since (10) represents a first-order stochastic process, there exists an equivalent representation of (10) of the form:

$$z_{t+1}^i = \phi(z_t^i) + e_{t+1}^i \quad (15)$$

where $\phi(\cdot)$ is an arbitrary smooth function.¹⁰ If $\phi(z_t^i) = \rho z_t^i$, we can write:

$$\begin{aligned} y_{t+1}^i &= X^i\theta + z_{t+1}^i \\ &= X^i\theta + \rho z_t^i + e_{t+1}^i \\ &= X^i\theta + \rho(y_t^i - X^i\theta) + e_{t+1}^i \end{aligned}$$

⁹To compute the long-term probability of a particular value of $y = Y$, we proceed as follows. Say we have 3000 values of $X^i\hat{\theta}$, each with frequency 1.

1. Outer loop: let $y = Y$.

- (a) Inner loop: Take a specific value of $X^i\hat{\theta}$. We have $f_W(X^i\hat{\theta}) = \frac{1}{3000}$.
- (b) Compute $\hat{z}^i = Y - X^i\hat{\theta}$.
- (c) Obtain $f_Z(Y - X^i\hat{\theta})$ using the ergodic distribution of \hat{z} . This is just the frequency of the discretized \hat{z}^i interval in which $Y - X^i\hat{\theta}$ happens to fall.
- (d) Repeat for all values of $X^i\hat{\theta}$ and take the sum of $f_Z(Y - X^i\hat{\theta})$ divided by 3000. This yields the probability that $y = Y$, which we have written $f_Y(Y)$. End of inner loop.

2. Repeat for all values Y to obtain all values of $f_Y(y)$.

Given that the algorithm is based on a discretization, we renormalize probabilities $f_Y(y)$ so that they exactly sum to 1.

¹⁰For our illustration, it is enough to assume that the errors e_{t+1}^i are not autocorrelated, but they need not be homoskedastic.

which can be rewritten:

$$y_{t+1}^i - y_t^i = (1 - \rho)X^i\theta + (\rho - 1)y_t^i + e_{t+1}^i \quad (16)$$

If y_t^i stands for log GDP per head, equation (16) is the standard conditional convergence model (Barro, 1991; Mankiw, Romer and Weil, 1992). What we estimate in this paper is a generalized version of model (16) where we replace the fixed parameter ρ with a smooth function $\phi(\cdot)$ to yield:

$$y_{t+1}^i = X^i\theta + \phi(y_t^i - X^i\theta) + e_{t+1}^i \quad (17)$$

The shape of function $\phi(\cdot)$ captures the way in which y_{t+1}^i converges to its steady state $X^i\theta$ or, in the Quah (1993) interpretation, how fast it returns to its conditional mean $X^i\theta$. Equation (16) can thus be seen as a generalization of the standard MRW model in which we do not impose linearity around the steady state and let the data tell us how rapidly the process converges depending on how far it is from its steady state. It can also identify the presence of multiple (deterministic) steady states and determine which ones are stable. As we will see, however, this approach to convergence is insufficiently informative if the true data generation process is stochastic because the shape of $\phi(\cdot)$ by itself tells us little about σ -convergence.

Function $\phi(\cdot)$ can be estimated by replacing, in equation (17), $y_t^i - X^i\theta$ with z_t^i (or a consistent estimate of it). After replacement, this boils down to applying a standard kernel regression to:

$$\hat{z}_{t+1}^i = \phi(\hat{z}_t^i) + e_{t+1}^i \quad (18)$$

In the unconditional case, we simply replace \hat{z}_t^i and \hat{z}_{t+1}^i with y_t^i and y_{t+1}^i .

3 The data

We now turn to the empirical implementation. As discussed in the introduction, our goal is to predict what the future distribution of economic activity over space would look like, should current tendencies persist. We use job figures as a proxy for economic activity. County-level sectoral employment data come from the Regional Economic Information System (REIS) compiled by the U.S. Bureau of Economic Analysis (BEA). We

use employment data for 1972 and 1992 in thirteen sectors, covering the entire economy: farming; agricultural services; mining; construction; manufacturing; transportation and utilities; wholesale; retail; FIRE (finance, insurance and real estate); other services; federal government; military; and state and local government. We focus on contiguous US because we believe that, over the period under investigation, labor and capital mobility towards Alaska and Hawaii were lower than now. Pooling them with the contiguous US may therefore not be appropriate for our purpose. We are left with 3092 counties. Sectoral employment data are missing for some counties, either because they are unavailable or because they are not disclosed.¹¹

To control for county-specific time-invariant characteristics, we use data on county area, latitude, and longitude from the U.S. Geological Survey (USGS). Counties are assumed to be centered at their county seat. The average county size is 2491 square kilometers, corresponding to an average diameter of approximately 50 kilometers (30 miles).¹² Counties vary considerably in size, however: the coefficient of variation of county area is 1.36. Western counties in particular tend to be larger than their eastern counterparts. Dummies are also created to control for whether a county is on a large body of water, such as a lake or ocean, or for whether it is on the border with Canada or Mexico. In particular, we include dummies for: the Atlantic ocean; the Pacific ocean; the Great Lakes; the gulf of Mexico; the Mexican border; and the Canadian border. Information of proximity to borders and water was compiled from detailed maps provided by the American Automobile Association (AAA). Latitude and longitude are also included as regressors. Finally, given that economic activity in the U.S. is concentrated on the Atlantic and the Pacific seaboards, we add dummies for counties located in states on the East coast or the West coast.

¹¹For some counties sectoral employment is not revealed in order not to violate employer confidentiality. For other counties sectoral employment is simply reported as ‘less than 10’; in those cases we set employment equal to 5.

¹²This approximation obviously underestimates the actual diameter, since counties are not perfect circles. It is nevertheless useful as a ballpark figure.

4 Results

4.1 σ -convergence and β -convergence

To get a feel for whether jobs have become more or less concentrated, *Table 1* reports the standard deviation of log employment at the county level in 1972 and 1992. A decreasing standard deviation reflects employment becoming more equally spread across counties — this is what Sala-i-Martin (1996) calls σ -convergence —; an increasing standard deviation points to employment becoming more concentrated in space. As can be seen, for total employment the tendency has been towards more concentration (divergence). At the sectoral level, there is a clear difference between service and non-service sectors. Services (‘retail’, ‘FIRE’, and ‘other services’) have become more concentrated; most other sectors, such as manufacturing and farming, have exhibited deconcentration (convergence).

Table 2 reports the results of a standard linear regression of employment growth on initial log employment — the standard test of unconditional β -convergence. A positive coefficient on initial employment points to concentration (divergence), whereas a negative coefficient indicates deconcentration (convergence). Our findings from *Table 1* are confirmed. We find concentration of employment at the aggregate level and in the service sectors, and deconcentration in most other sectors. This suggests that services are driving aggregate employment dynamics. This should not come as a surprise, given the weight of services in the economy: in 1972 ‘retail’, ‘FIRE’ and ‘other services’ already made up 42% of total employment; by 1992 this share had grown to 54%.

The increasing concentration of aggregate county employment stands in contrast with the observed employment deconcentration across metropolitan areas (Carlino and Chatterjee, 2001). This suggests deconcentration across large counties, and concentration across smaller sized counties. A quick-and-easy way of checking this is to split up our sample into two groups: the 200 counties with more than 86,000 workers in 1972, and the remaining 2875 counties. As expected, for the group of large counties we find deconcentration across the board. *Table 3* shows a negative coefficient on initial log employment for all sectors. In contrast, for the group of small counties *Table 4* shows total

employment becoming more concentrated. Again, there is a dichotomy across sectors: concentration in services, and deconcentration in the rest of the economy. Summing up, standard convergence analysis indicates that the tendency towards employment deconcentration only holds for a limited group of high employment counties. For the rest of the distribution, concentration seems to be the norm.

4.2 Unconditional kernel regressions

Splitting up the sample into two parts, and running linear regressions on each part, is a rather rudimentary way of dealing with nonlinearities. A more appropriate way of capturing the richness of the dynamics is to run nonlinear kernel regressions on the entire sample. The unconditional estimating equation is of the form:

$$y_{t+1}^i = \phi(y_t^i) + e_{t+1}^i \quad (19)$$

where employment data have been normalized, assuming an exponential trend.¹³ To facilitate interpretation, we have plotted the employment *growth* as a function of initial employment.¹⁴ In this case, a negative slope indicates dispersion (convergence) and a positive slope indicates concentration (divergence). Also, when the curve cuts the horizontal axis from above, we have a stable equilibrium if the underlying model is deterministic; when it cuts it from below, we have an unstable equilibrium.

Figure 1 plots the results for total employment and by sector. We start by looking at the picture for total employment ('Total'). The curve cuts the horizontal axis three times: in a deterministic model, the two equilibria at the extremes would be stable and the middle equilibrium unstable. This suggests that at intermediate values of initial employment, forces exist that push total employment away towards the extremes. In other words, the middle part of the distribution exhibits divergence in the deterministic sense: if employment starts off below the middle equilibrium, the county is on average predicted to lose jobs, and is expected to converge towards the low steady state. In contrast, if

¹³We also experimented with a linear trend; this did not significantly affect the quantitative results though.

¹⁴This amounts to subtracting y_{72}^i from both sides of equation (19).

employment starts off above the middle equilibrium, the county is expected to gain jobs and converge towards the high steady state. This result illustrates how misleading standard linear regressions can be when testing β -convergence. Kernel regressions represent a significant improvement over standard linear tests of convergence.

It turns out that exactly 50% of counties,¹⁵ i.e., all counties with less than 7664 jobs in 1972, are predicted to slowly empty out according to the deterministic model; the remaining 50% of counties are predicted to gain jobs and to end up in the high steady state. We only observe convergence in the lower and the upper end of the distribution, where the slope is negative. Regarding the upper tail, the 9% biggest counties – corresponding to those with more than 57,611 jobs in 1982 – are predicted to lose jobs on average. These results are by and large consistent with the results from the previous sub-section.

Turning to individual sectors, we see that for most of the non-service sectors – such as ‘manufacturing’ and ‘farming’ – the slope is monotone and decreasing. This suggests deconcentration (convergence). The deterministic steady state is where the curve cuts the horizontal axis. In contrast, for two out of the three service sectors – ‘retail’ and ‘FIRE’ – the picture resembles that of aggregate employment. In ‘retail’, for instance, there is a steady state with low retail employment and a steady state with high retail employment, with the middle group disappearing. More specifically, the deterministic model predicts that 48% of counties will end up in the low steady state (with 123 retail jobs, using 1972 figures) and the remaining 52% will end up in the high steady state (with 16,874 retail, using 1972 figures).

4.3 Unconditional ergodic distributions

The kernel regressions are based on a deterministic underlying model, with all counties ending up in a small number of steady states. Predictions about future concentration are therefore extremely crude. To generate more realistic predictions, we turn to a stochastic model in which the long-term behavior of the system is characterized not by one or two

¹⁵To be precise, 1529 out of 3069 counties.

steady states but by an entire distribution – the ergodic distribution.

Figure 2 plots the ergodic distributions for all sectors, and compares them to the distributions of employment in 1972 and 1992. The method used to derive the ergodic distribution was described in the methodology section. If the ergodic distribution is entirely below that in 1972 and 1992, this suggests deconcentration: based on current trends, counties are predicted to look more alike in the future than in the past. If the ergodic distribution has a thinner upper tail than the actual employment distributions for 1972 and 1992, this means fewer high employment counties in the future. In contrast, if the ergodic distribution has a 'hat shape' instead of a 'bell shape', with more mass on high employment values, there will be more concentration in the future: the number of high employment counties is predicted to increase while the number of intermediate employment counties is expected to fall.

It is important to realize that these predictions about the long-term distribution of employment across counties assume that transition probabilities will remain constant. This need not be true, as changes in demand and technology may modify them. The ergodic distribution is therefore best understood as an indicator of current long-term tendencies.

Results shown in *Figure 2* fall basically into two categories: 'total', 'retail', 'FIRE' and 'military' employment exhibit a hat shape, while all other sector exhibit a (slightly skewed) bell shape. For the case of 'FIRE' we even get 'twin peaks' – a bimodal distribution. Results for two sectors — 'agricultural services' and 'mining' — stand apart from the rest, because they are dominated by a very large number of counties with zero or undisclosed employment.

These results by and large confirm our findings in the kernel regressions. Total employment is expected to become more concentrated over time. This means that US counties will tend to become more differentiated in terms of employment, with more counties with little if any employment, more counties with high employment, and fewer with intermediate employment. This phenomenon at the aggregate level is a reflection of increased concentration in 'retail' and 'FIRE', and to a much smaller extent — the sector

is small — in ‘military’.

All other sectors are predicted to either remain at their current level of concentration or to become less concentrated. We indeed see that, for most of these sectors, the ergodic distribution is tighter than the current distribution. Furthermore, we also see that, for ‘farming’, ‘manufacturing’, ‘transportation and utilities’, and ‘wholesale’, the ergodic distributions shift to the left compared to past distributions. This suggests fewer counties with large employment levels and more counties with low or intermediate employment levels. For these sectors, convergence is predicted to take place at levels of county employment lower than current levels. This is consistent with a spreading of these activities over space. In contrast, concentration in ‘construction’, ‘federal civilian’ employment, ‘state and local’ administration, and ‘other services’ is predicted to remain basically constant. In the case of ‘federal employment’ and ‘other services’ there remains a fairly thick upper tail though, suggesting the continued presence of a small number of counties with large employment in those sectors.

If we compare *Figure 2* with *Figure 1*, we note that hat-shaped ergodic distributions arise whenever the kernel regression has two clearly identified stable equilibria. This is not surprising: the two deterministic steady states can be thought of as attracting points to which employment realizations tend. Stochastic shocks, however, ensure that employment does not settle in either steady state. Intuitively, the hat shape of the ergodic distribution results from the ‘mixing’ of the distributions around each of the two deterministic steady states. The kernel regression, however, fails to predict the extent of mixing and is therefore less informative.

In contrast, whenever the kernel regressions suggest convergence to a single steady state, the ergodic distribution exhibits a bell shape. The kernel regression does not, however, indicate whether future concentration will differ from current one. This information is only obtained by calculating the ergodic distribution. Finally, when the kernel regression has a flat – albeit negative – segment above the steady state, as in the case of federal employment and other services, we observe that the ergodic distribution has a fat upper tail. The size of this upper tail cannot, however, be assessed solely on the

basis of the shape of the kernel regression. We therefore see that, in all cases, kernel regressions — which are themselves a generalization of standard β -convergence tests — are less informative than the ergodic distribution in identifying the predicted direction of change.

4.4 Conditional kernel regressions and ergodic distributions

The analysis presented so far may be misleading if the distribution of employment across counties partly reflect time-invariant differences. Ergodic distributions computed without conditioning on these differences may underestimate the magnitude of stochastic shocks and thus misrepresent the long-term distribution of employment across counties.

To deal with this problem, we turn to the conditional model discussed in Section 2. As explained there, we first run a pooled regression of county employment on county characteristics. This regression has the form (12). The X^i characteristics include a variety of geographical features for which we have data. We then obtain the \hat{z}_t^i using (13). These \hat{z}_t^i are then used to calculate a new set of kernel regressions and a new set of ergodic distributions using formula (14).

Figure 3 plots the outcome of the kernel regressions of the \hat{z}_t^i obtained using (18). This kernel regression measures whether the stochastic component of y_t^i has a deterministic steady state. As explained in Section 2, this approach is equivalent to the standard conditional β -convergence approach, except that it allows for non-linearities.

Comparing *Figure 3* with *Figure 1*, we again get divergence of the middle and convergence in the lower tail for ‘total’ employment, though there is less evidence of convergence in the upper tail. Turning to individual sectors, we again see that the non-service sectors exhibit deconcentration (convergence), whereas two out of the three service sectors — ‘retail’ and ‘FIRE’ — and to some extent the military are predicted to become more concentrated over time. In other words, conditioning on time-invariant county characteristics does not change the basic story: there is concentration at the aggregate level, and this concentration is driven by the service sectors. This implies that observed concentration is not due to geographical differences between counties.

Conditional ergodic distribution are presented in *Figure 4*. The basic pattern is similar to that depicted in *Figure 2*: ‘total’ employment and ‘retail’ exhibit a hat shaped ergodic distribution, suggesting further concentration in the future. But once we condition on geographical characteristics, ‘FIRE’ and the ‘military’ no longer have a hat shape, although their ergodic distributions continue to display a higher variance and fatter upper tails than past distributions. For ‘state and local’ administration, employment is expected to become more concentrated over time once we condition on geographical factors. Taken together, these results suggest increasing concentration in a number of service sectors.

Other results are broadly similar to *Figure 2*, although in several instances the shift to the left, indicating deconcentration, is less noticeable once we condition on geographical characteristics. This is especially true for ‘farming’. This is hardly a surprise given that geographical characteristics are thought to play an important role in the location of agricultural activity. The tendency towards less concentration for ‘manufacturing’ is confirmed, although the conditional ergodic distribution exhibits a larger variance than the unconditional one, again suggesting that geographical factors play a role in the positioning of manufacturing enterprises.

5 Conclusion

In this paper we have examined how the distribution of employment across US counties is likely to evolve if current concentration and deconcentration forces remain unchanged. To this effect, we developed a methodology borrowing heavily from the work of Quah and building upon the literature on β and σ -convergence. We computed non-parametric β -convergence regressions, conditional and unconditional. Using non-parametric methods, we also computed detailed ergodic distributions for total employment and sectoral employment across US counties.

If the forces that currently affect the spatial concentration of economic activity continue to operate in the same direction, our results suggest that employment will become increasingly concentrated across counties. Although very large cities may lose jobs,

the proportion of counties with modal employment is expected to decrease in favor of medium to high employment level counties. More specifically, the 9% largest counties and the 15% smallest counties exhibit deconcentration; the remaining 76% exhibit concentration. This result is consistent with deconcentration across urban areas (Carlino and Chatterjee, 2001) and concentration across US counties (Desmet and Fafchamps, 2003). Of course, given that the 9% high employment counties accounted for two thirds of total employment in 1972, the overall picture is a complex one, in which the importance of these counties is reduced due to deconcentration, but the remaining counties become more differentiated.

There are important differences across sectors. As in the rest of the literature, we find deconcentration in manufacturing. Deconcentration is also the norm in other non-service sectors. However, a number of service activities are becoming more concentrated, namely ‘retail’ trade and ‘finance, insurance and real estate’. Given the importance of these sectors, they drive the evolution of the spatial distribution of total employment. Limiting the focus of analysis to manufacturing is misleading. The US is a service economy, and services are behaving very differently from the other sectors.

On the methodological side, our research shows the importance of using non-parametric methods and of looking at the entire distribution, not just at cities, when studying convergence. It also demonstrates that, even when β -convergence tests are conducted in a non-parametric way – as done here – they are not sufficiently informative. Computing the ergodic distribution associated with a given set of transition probabilities is a more useful way of analyzing the likely evolution of concentration in the future. With the approach presented here, it is also straightforward to condition on time-invariant characteristics in a way that is fully consistent with standard analysis of conditional β -convergence. The methodology developed here can easily be applied to the study of any distributional dynamics.

This paper leaves a number of questions unanswered. First, it is unclear whether the forces identified here operate in a similar manner in other time periods and other parts of the world. Applying the same approach to other data sets is necessary before

we can conclude that the process described here generalizes beyond the confines of this study. Second, the methodology presented here does not (yet) allow statistical inference in the normal sense. Statistical tests are reported for some of the statistics presented here – notably confidence intervals for kernel regressions and t -statistics for the coefficients of time-invariant regressors in the conditioning regressions (12). But we do not present a ‘test’ of (conditional or unconditional) convergence based on estimated ergodic distributions. In principle, such a test could be developed provided an intuitively satisfying counter-factual distribution could be devised. It should also be possible to use bootstrapping to test whether the mode of the ergodic distribution has shifted to the left or the right relative to the current distribution (Kremer, Onatski and Stock, 2000). Developing such tests is left for future research.

References

- [1] Beeson, P., DeJong, D., and Troesken, W., 2001. Population Growth in US Counties, 1840-1990, *Regional Science and Urban Economics*, 31(6), 669-699.
- [2] Barro, R.J., 1991. Economic Growth in a Cross-Section of Countries, *Quarterly Journal of Economics*, 106(2), 407-443.
- [3] Blanchard, O. and Katz, L., 1991. Regional Evolutions, *Brookings Papers on Economic Activity*, 1, 1-75.
- [4] Carlino, G.A. and Chatterjee, S., 2001. Aggregate Metropolitan Employment Growth and the Deconcentration of Metropolitan Employment, *Journal of Monetary Economics*, 48, 549-583.
- [5] Desmet, K. and Fafchamps, M., What Are Falling Transport Costs Doing to Spatial Concentration across U.S. Counties?, *CEPR Discussion Paper* 3853.
- [6] Dumais, G., Ellison, G. and Glaeser, E., 2002. Geographic Concentration as a Dynamic Process, *Review of Economics and Statistics*, 84(2), 193-204.

- [7] Durlauf, S.N. and Quah, D.T., 1999. The New Empirics of Economic Growth, in: J.B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, volume 1A, Chapter 4, North Holland Elsevier Science: Amsterdam.
- [8] Eaton, J. and Eckstein, Z., 1997. Cities and Growth: Theory and Evidence from France and Japan, *Regional Science and Urban Economics*, 27, 443-474.
- [9] Fujita, M., Krugman, P. and Venables, A.J., 1999. *The Spatial Economy: Cities, Regions, and International Trade*, MIT Press: Cambridge, Mass.
- [10] Henderson, J.V., 1988. *Urban Development: Theory, Fact and Illusion*, Oxford University Press: New York.
- [11] Higgins, M., Levy, D. and Young, A., 2003. Growth and Convergence across the US: Evidence from County-Level Data, Emory Economics Working Paper 0306.
- [12] Kim, S., 1995. Expansion of Markets and the Geographic Distribution of Economic Activities: The Trends in U.S. Regional Manufacturing Structure, 1860-1987, *Quarterly Journal of Economics*, 110, 881-908.
- [13] Kolko, J., 1999. Can I Get Some Service Here?, mimeo, Harvard University.
- [14] Kremer, M., Onatski, A. and J. Stock, 2000. Searching for Prosperity, Department of Economics, Harvard university, mimeo.
- [15] Luenberger, D.G., 1979. *Introduction to Dynamic Systems*, Wiley: New York.
- [16] Mankiw, N.G., Romer, D. and Weil, D.N., 1992. A Contribution to the Empirics of Economic Growth, *Quarterly Journal of Economics*, 107, 407-437.
- [17] Mood, A.M., Graybill, F.A., and Boes, D.C., 1974. *Introduction to the Theory of Statistics*, McGraw Hill.
- [18] Quah, D.T., 1993. Galton's Fallacy and Tests of the Convergence Hypothesis, *Scandinavian Journal of Economics*, 95(4), 427-443.

- [19] Quah, D.T., 1996a. Empirics for Economic Growth and Convergence, *European Economic Review*, 40, 1353-75.
- [20] Quah, D.T., 1996b. Convergence Empirics Across Economies with (Some) Capital Mobility, *Journal of Economic Growth*, 1, 95-124.
- [21] Quah, D.T., 1997. Empirics for Growth and Distribution: Stratification, Polarization, and Convergence Clubs, *Journal of Economic Growth*, 2, 27-59.
- [22] Sala-i-Martin. X., 1996. Regional Cohesion: Evidence and Theories of Regional Growth and Convergence, *European Economic Review*, 40, 1325-52.
- [23] Stokey, N.L. Lucas, R.E. and Prescott, E.C., 1989. *Recursive Methods in Economic Dynamics*, Harvard University Press: Cambridge, Mass.

Table 1: Standard deviations of sectoral employment in 1972 and 1992 in logs

Sector	standard deviation (logs) 1972	standard deviation (logs) 1992
Total	1.34	1.43
Farming	.94	.84
Agricultural Services	1.32	1.31
Mining	1.67	1.78
Construction	1.51	1.58
Manufacturing	1.97	1.88
Transportation/Utilities	1.54	1.57
Wholesale	1.73	1.67
Retail	1.42	1.58
FIRE	1.54	1.64
Other Services	1.50	1.62
Federal Civilian	1.57	1.60
Military	1.49	1.50
State/Local	1.31	1.36

Source: REIS, Bureau of Economic Analysis

Table 2: Sectoral employment growth on initial sectoral employment (all 3074 counties)

Dependent variable: annual growth rate in sectoral employment 1972-92							
	Total	Farming	Ag serv	Mining	Constr	Manuf	Trans/util
const	-.0045894	.031724	.0795073	.0319508	.0232263	.054432	.0356632
	-2.346	18.468	37.606	12.668	11.562	26.750	18.522
emp72	.0022102	-.006113	-.0072239	-.0054879	-.0009096	-.0060793	-.0030251
	10.385	-24.146	-15.188	-9.313	-2.805	-21.744	-9.506
R^2	0.0399	0.1595	0.0855	0.0368	0.0027	0.1372	0.0305
	Wholesale	Retail	FIRE	Other serv	Fed civ	Milit	State/Loc
const	.0554855	-.0042489	.0031527	.0172376	.009566	.0058711	.0127573
	28.480	-2.294	1.950	8.878	7.904	5.114	8.322
emp72	-.0043428	.0032981	.0019356	.0021457	-.0007685	-.0011588	.0005706
	-12.543	12.969	7.445	8.147	-3.222	-5.269	2.660
R^2	0.0528	0.0521	0.0184	0.0221	0.0034	0.0090	0.0023

Absolute values of t-statistics (corresponding to spatially corrected standard errors) in brackets.

Table 3: Sectoral employment growth on initial sectoral employment (200 largest counties)

Dependent variable: annual growth rate in sectoral employment 1972-92							
	Total	Farming	Ag serv	Mining	Constr	Manuf	Trans/util
const	.0654953	-.0020957	.0979573	.0842865	.081878	.1346491	.0934903
	3.960	-0.538	9.623	7.272	4.892	8.551	6.686
emp72	-.0037547	-.0010493	-.0071837	-.0107865	-.0075326	-.0130225	-.0084311
	-2.785	-1.901	-4.790	-5.568	-4.195	-8.736	-5.605
R^2	0.0377	0.0179	0.1083	0.1409	0.0820	.2782	0.1375
	Wholesale	Retail	FIRE	Other serv	Fed civ	Milit	State/Loc
const	.1189618	.0997601	.0465353	.0795902	.0362754	.0362712	.0364982
	8.686	6.481	3.841	6.284	4.601	4.557	3.155
emp72	-.0104643	-.0074854	-.0023342	-.0035878	-.003329	-.004913	-.0020426
	-7.031	-5.065	-1.855	-3.012	-3.621	-5.167	-1.777
R^2	0.1998	0.1147	0.0171	0.0483	0.0621	0.1188	0.0157

Absolute values of t-statistics (corresponding to spatially corrected standard errors) in brackets.

Table 4: Sectoral employment growth on initial sectoral employment (2874 smallest counties)

Dependent variable: annual growth rate in sectoral employment 1972-92							
	Total	Farming	Ag serv	Mining	Constr	Manuf	Trans/util
const	-.0116011	.0403319	.0898732	.0330733	.0198499	.0569191	.0414091
	-4.686	21.371	37.196	12.556	8.166	25.047	17.571
emp72	.0030247	-.0074054	-.0101888	-.006341	-.0002885	-.0065219	-.0041513
	10.924	-26.547	-17.670	-9.845	-0.701	-20.040	-10.105
R^2	0.0399	0.1967	0.1202	0.0440	0.0002	0.1265	0.0368
	Wholesale	Retail	FIRE	Other serv	Fed civ	Milit	State/Loc
const	.0597689	-.0156651	.0045739	.0176736	.0123751	.0043236	.0089734
	25.129	-6.931	2.254	7.262	8.304	3.144	4.688
emp72	-.0052954	.0050313	.0016601	.0020689	-.0014565	-.0008204	.0011516
	-11.703	15.613	4.827	6.015	-4.630	-2.939	4.150
R^2	0.0496	0.0785	0.084	0.0130	0.0074	0.0030	0.0060

Absolute values of t-statistics (corresponding to spatially corrected standard errors) in brackets.

Figure 1: unconditional kernel regressions









Figure 1: unconditional kernel regressions (cont'd)

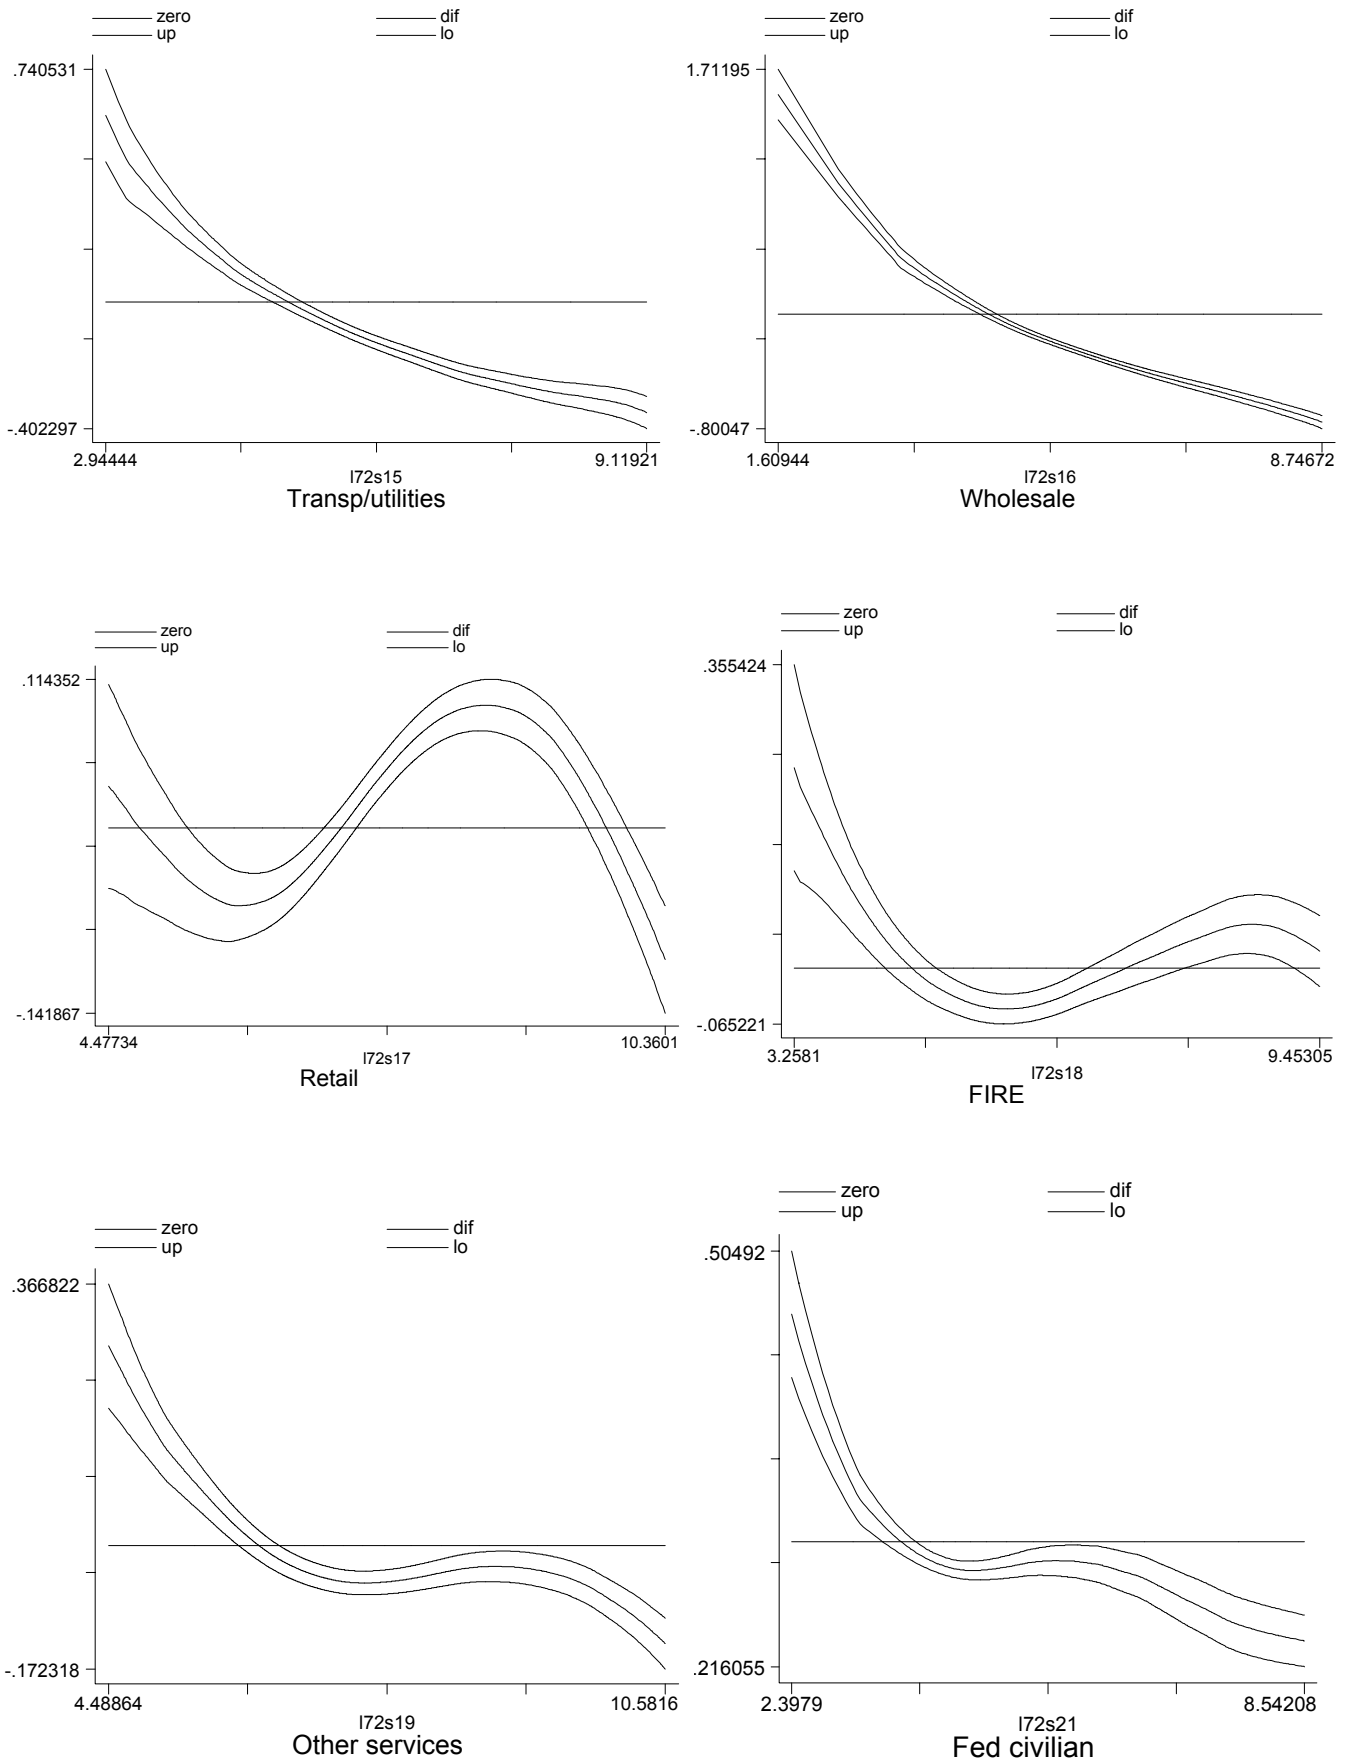


Figure 1: unconditional kernel regressions (cont'd)





Figure 2: Unconditional ergodic distributions









Figure 2: Unconditional ergodic distributions (cont'd)

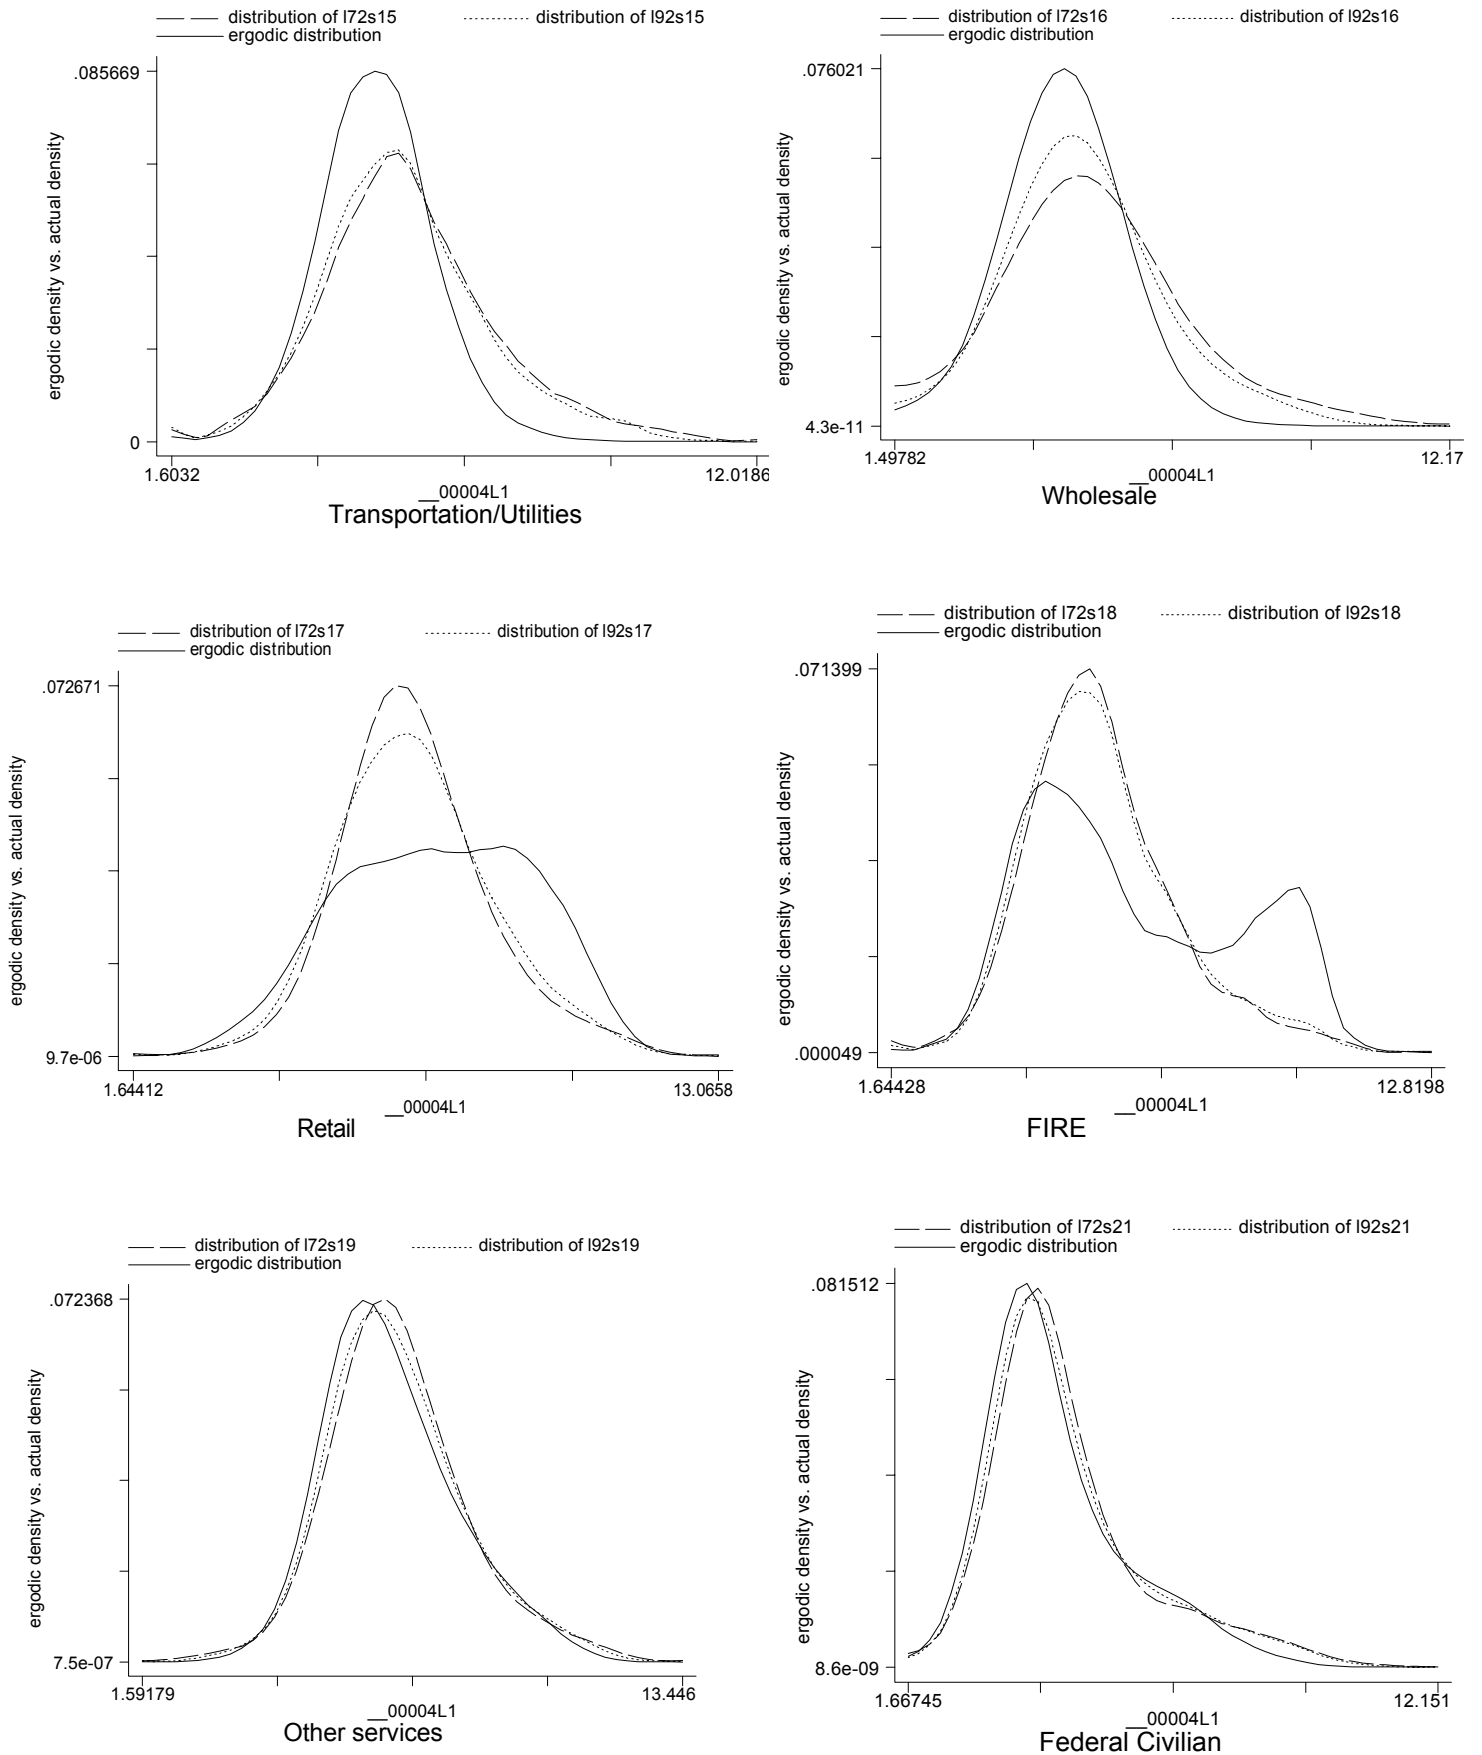


Figure 2: Unconditional ergodic distributions (cont'd)





Figure 3: conditional kernel regressions









Figure 3: conditional kernel regressions (cont'd)









Figure 3: conditional kernel regressions (cont'd)





Figure 4: conditional ergodic distributions

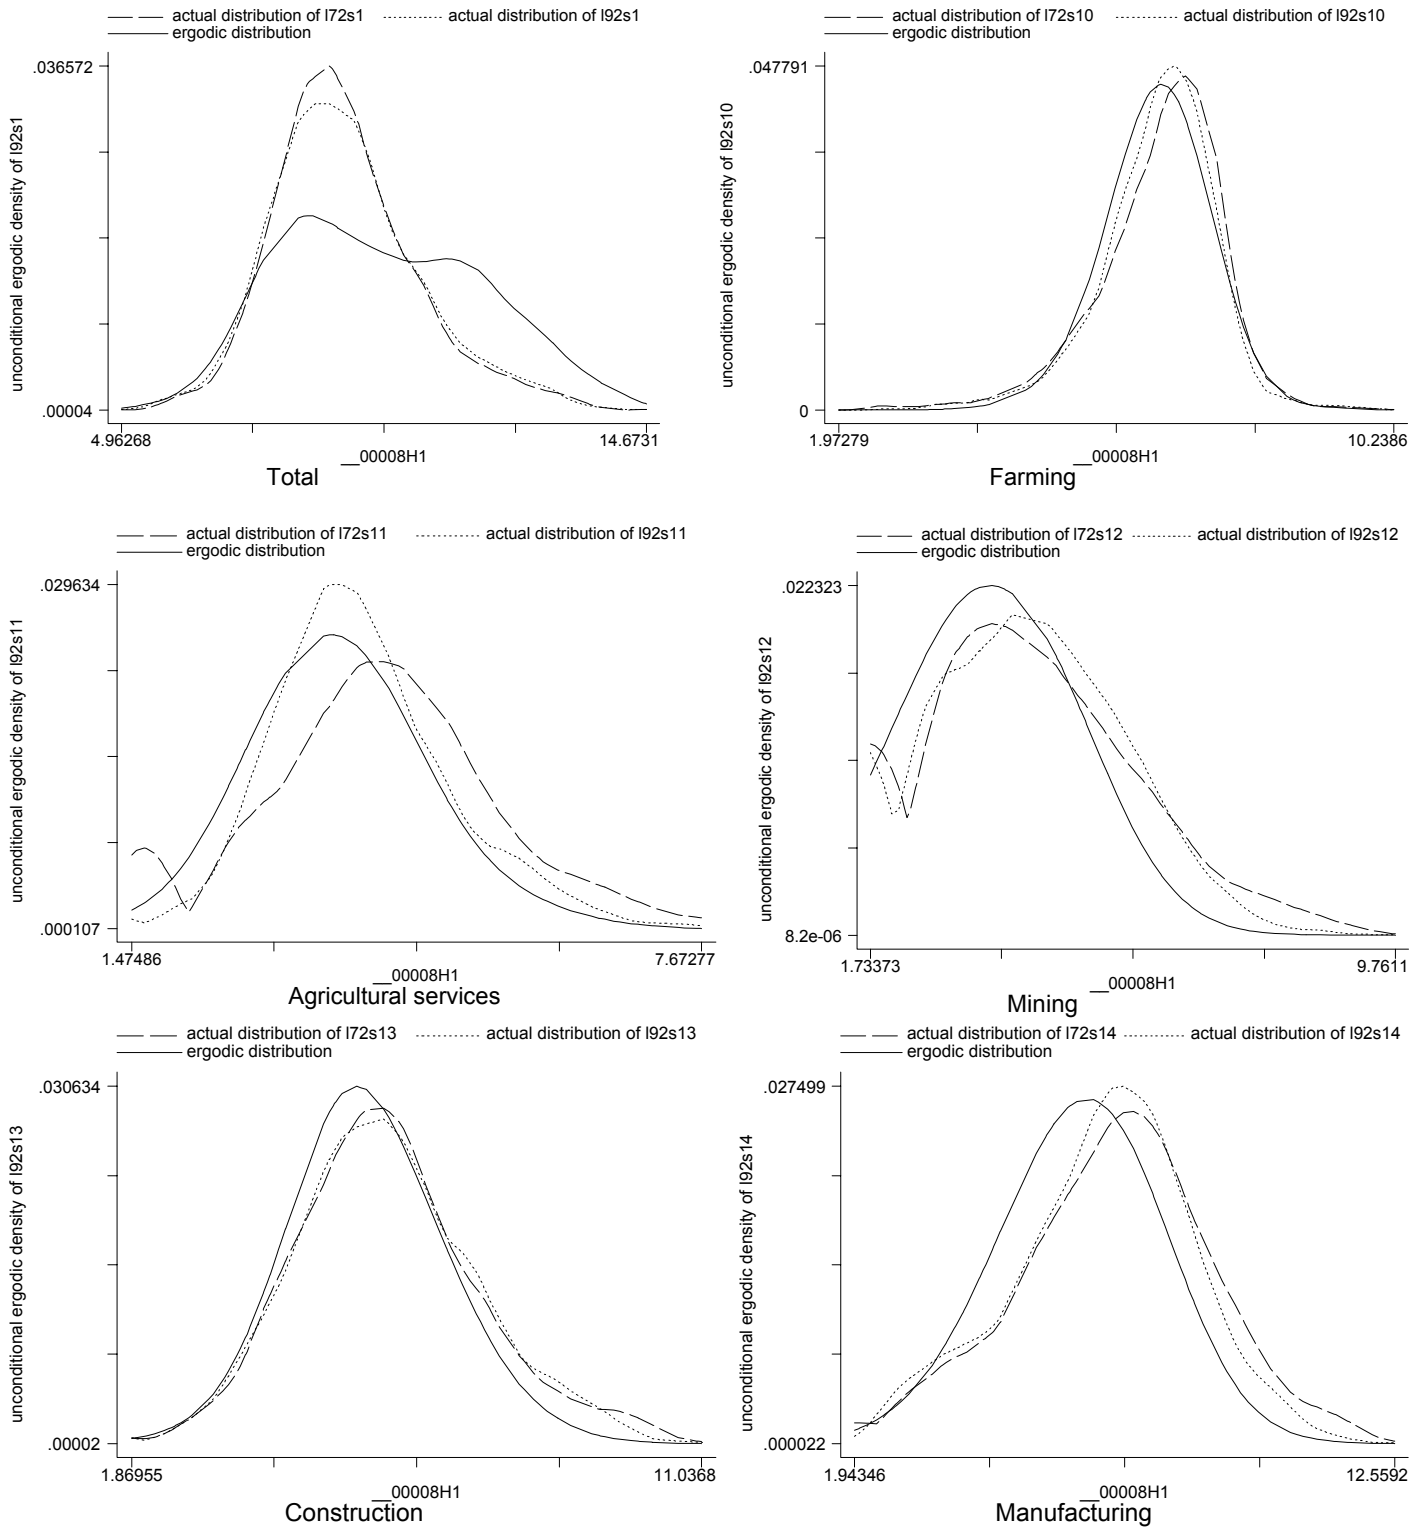


Figure 4: conditional ergodic distributions (cont'd)

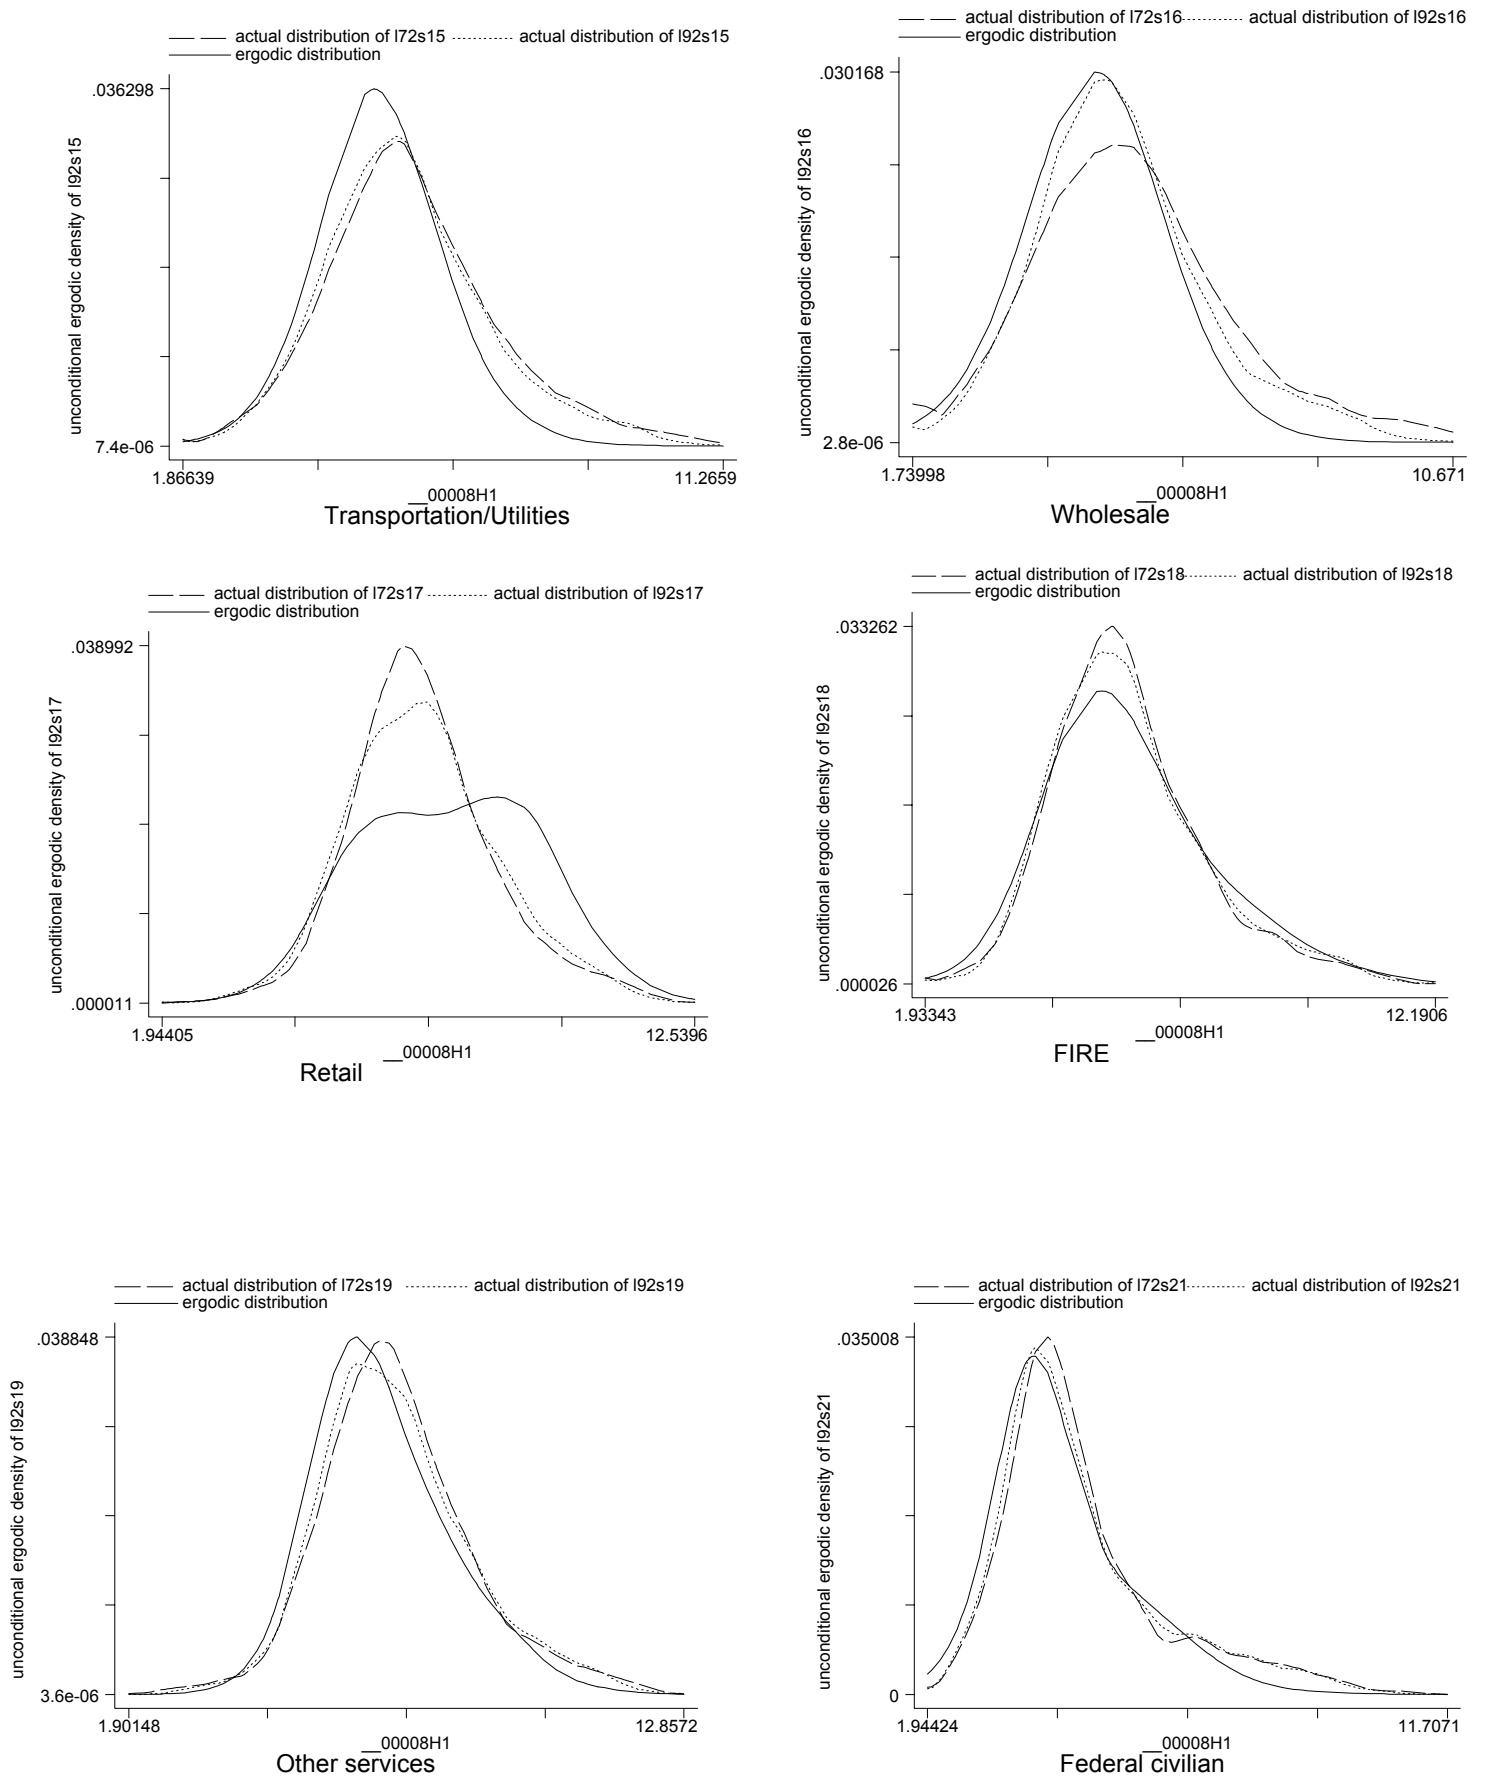


Figure 4: conditional ergodic distributions (cont'd)

