Journal of Monetary Economics 000 (2017) 1–4

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmoneco

Comment on "Tarnishing the golden and empire states: Land-use regulations and the U.S. economic slowdown," by Herkenhoff, Ohanian, and Prescott

Christopher Tonetti*

Stanford University, Graduate School of Business, Stanford, CA 94305-7298, United States

A R T I C L E I N F O

Article history:
Received 1 November 2017
Accepted 2 November 2017
Available online xxx

1. Introduction

Over the past few decades migration across U.S. regions has declined, house prices in the most productive regions have risen sharply, and the growth rate of aggregate output per worker has slowed. Land is a crucial immobile factor of production in housing and in the production of goods and services and, as such, land-use is a natural candidate for explaining the comovement in migration, housing, and output. Taking the perspective that stricter land-use regulations decrease the effective amount of available land, this paper first documents how land-use policy has changed over time and across regions and then computes the impact of these changes on output per worker.

The task of computing the impact of land-use policy is complicated by the fact that there are no direct measures of such land-use policies. Much of the literature that studies land-use regulation uses a single 2008 cross-section of the Wharton Land Regulation Index (WLRI), which omits the ability to make any statements about changes over time. The key innovation in Herkenhoff et al. (2018) (hereafter HOP) is to treat land-use regulation as a latent variable and filter the time series using equilibrium relationships from a structural model. This brazen exercise allows the paper to answer a policy-relevant question of central importance to the welfare of the U.S. population, instead of resigning the economic profession to silence because the necessary data was neither collected historically nor at scale. According to the model, increased restrictions in land use have caused a spatial misallocation of people and capital, resulting in a substantial decrease in the level of output per worker today.

This comment will first review the methods used to construct the measure of land-use regulation. It will then briefly compare the model used in this paper to those prominent in the literature and discuss the source of differing predictions. Then, it will discuss the counterfactual exercise, with a focus on alternative interpretations of the causal relationship between land-use policy and other key determinants of regional population and output. Finally, it will offer some concluding remarks with an eye toward future research.

* Corresponding author.
E-mail address: tonetti@stanford.edu

https://doi.org/10.1016/j.jmoneco.2017.11.002

2. Constructing a time series for regional land-use regulations

While this is not the first paper to study the impact of land-use regulations on output, it stands out in its ability to study the change in land-use policy over time and by its scale, at the level of the U.S. states and aggregate economy.[1] Since such land-use regulations are not directly measured over time, the structural model used for filtering is at the center of the exercise.[2] The dynamic general equilibrium model features land as a fixed factor of production and allows changes in the aggregate capital stock and the regional allocation of capital and people to offset land-destroying regulations. The logic of the model is simple: send people where productivity is high, where effective land is plentiful, and where local amenities are high quality. The difficulty is in separately identifying unmeasured local amenities, productivity, and the effective amount of land. To gain identification, the authors first construct a state-level panel dataset of real output per worker, population, real house prices, and urban land acreage by weaving together data from the BLS, Census, BEA, Census of Housing, American Community Survey, and USDA Economic Research Services. This data, together with Cobb-Douglas production functions for housing and consumption goods, and a housing requirement linking housing quantity to population, allows the panel of amenities, land-use regulations, and productivity to be identified, as well as the amount of capital and land allocated to the consumption good and housing sectors in each region-time.[3] Using the recovered parameters and latent variables, the authors perform various counterfactuals that demonstrate that the tightening of land-use regulations caused large decreases in output per capita in the U.S., differentially across regions.

3. How costly are land-use regulations?

In the neoclassical HOP economy, if in 2014 all states were deregulated back to their 1980 levels, steady state aggregate output would be 7.2% higher; similarly deregulating just California and New York would raise aggregate output by 3.7%. The family that represents the whole U.S. population would require 7.3% higher consumption every year as compensation in order to be willing to choose to live in the world with 2014-level regulations instead of 1980 regulations. These numbers are large, and point to land-use regulation reform as a promising channel to improve the well-being of U.S. residents. These predictions, however, are much smaller than some other papers. To make this comparison concrete, I will use Hsieh and Moretti (2017) (hereafter HM) as a foil, since I believe it is the closest paper to address a similar question at similar scope. HM report that the annual output growth rate between 1964 and 2009 would have doubled from 0.8% to 1.6% if house prices were held fixed at their 1964 levels.[4] With a focus on a very similar policy question, why is there such a difference in predicted outcomes?

3.1. Level versus growth effects

First, to make the comparison, steady state counterfactual levels from the HOP analysis must be converted to growth rates. The paper treats data from each year as that generated in the steady state equilibrium of a neoclassical growth model. It computes counterfactuals for a given year by computing the steady state for that year with a different land-use policy parameter α. In the neoclassical growth model, changes in α induce level changes in output, but not changes in output growth. Thus, all growth rates reported in the paper take a given change in the level of output across steady states and spreads that level change over the number of years between the initial and final periods of comparison. Although not directly comparable, the closest calculation to that in Hsieh and Moretti (2017) that is available in the HOP paper is to compare the counterfactual growth rate of aggregate output between 1980 and 2014 if in 1980 the U.S. had frozen its land-use regulations, keeping amenities and productivity in 2014 as measured.

Let Y_{t_0} be aggregate output in initial year t_0, T be the final year, and g be the measured average annual growth rate in output between t_0 and T. Let g_c be the counterfactual computed growth rate between t_0 and T in response to a counterfactual land-use policy that changes the level of steady state aggregate output in year T by a factor of c. Then,

$$Y_{t_0} e^{g(T-t_0)} = cY_{t_0} e^{g_c(T-t_0)} \quad (1)$$

$$g_c - g = \frac{\ln(c)}{T - t_0} \quad (2)$$

As reported in Column 7 of Table 6, counterfactual steady state output in 2014 would be 7.2% higher (i.e., c=1.072) if regulations in 2014 were set at their 1980 level. HOP report that spreading this increase in the level of output of 7.2% from 2000 to 2014 would increase annual output growth by 0.5 percentage points. Because land-use policies induce level

[1] Previous papers (e.g., Glaeser et al., 2005a; 2005b; Glaeser and Ward, 2009) use methods and data well-suited to cities, but rely on data not available at higher levels of aggregation needed for a macro analysis.

[2] Even if land-use policies were measured over time, there would still be a need for mapping the myriad measured policies into model variables. There is no model-free analogue of such a diverse collection of land-use policies in economically meaningful units.

[3] The authors use house price data, while the model features a rental market for housing. See Kiyotaki et al. (2011) for some implications of this way of matching theory and data.

[4] It is not clear why Hsieh and Moretti (2017) have an annual growth rate of output of 0.8%, when NIPA numbers over these decades would be closer to 3% or 2% per capita.

Please cite this article as: C. Tonetti, Comment on "Tarnishing the golden and empire states: Land-use regulations and the U.S. economic slowdown," by Herkenhoff, Ohanian, and Prescott, Journal of Monetary Economics (2017), https://doi.org/10.1016/j.jmoneco.2017.11.002

effects, spreading this same 7.2% level increase over 1980–2014 would generate an increase in annual output growth of 0.2 percentage points.

Connecting to Hsieh and Moretti (2017) by freezing land-use regulations all the way back to 1964 levels would have two countervailing forces. First, spreading the given 1980 counterfactual level increase in output of 7.2% over more years would reduce the increase in the annual growth rate to 0.15 percentage points per year. Second, fixing land-use policies at 1964 levels would increase counterfactual 2014 output relative to holding land-use policies at 1980s levels—much of the decrease in α occurred during the 1960s and 1970s. According to HOP calculations generated for this comment, output in 2014 would be 16.4% higher if land-use policy was kept at its 1960 level. Spreading this increase over 54 years results in an increase in output growth of 0.28 percentage points per year, greatly improving welfare by an 18 percent consumption equivalent measure, but much smaller than the 0.8 percentage points per year or doubling of growth as found in Hsieh and Moretti (2017).

3.2. Comparing key model features across papers

Why do HOP predict a large, but much smaller, cost of land-use regulation than Hsieh and Moretti (2017). There are four features of their macro analysis that differ from the standard approach in the more applied-micro tradition. The main differences are around aggregate labor supply, housing supply, the mechanism that allocates people to regions, and the economic measure of land-use regulations.

First, as a GE macro model, the level of the capital stock and the amount of labor supplied are endogenous. When the effective amount of land is scarce, other factors of production can be supplied to offset an aggregate loss in output. HM allow for the capital stock to change in response to land-use regulation in a manner similar to HOP, but HM normalize aggregate labor to 1 in all cases, and do not allow for changes in labor force participation distinct from changes in population. This labor supply effect, however, is likely to be a quantitatively small contributor to the difference in the predicted impact of land-use regulations, as the aggregate labor supply changes in the HOP counterfactuals are small. For example, deregulating to the 1960 level of α results in less than a 1% difference in aggregate labor supply.

Second, and probably the most important difference across models, is that in HOP land is a factor in the production of housing and housing supply equals housing demand, with house prices clearing the market. In HM, land-use regulations are modeled as a parameter that is the inverse elasticity of housing supply to population. In HOP, when effective land is in short supply, perhaps due to regulations, capital can be substituted for land to increase housing supply. This means that the housing stock need not contract as severely in high productivity regions in response to a tightening of land-use regulations, mitigating the negative impact of land-use regulation on house prices, and thus the spatial misallocation of people and capital. Put simply, this difference in modeling delivers a different elasticity of housing prices to land-use regulations, which can have a large impact on the elasticity of output to such regulations.

A feature shared by both types of models is that in response to more restrictive land-use policies, capital and labor can be substituted for land to buoy output. Since capital and labor are mobile and allocated optimally across regions, even when holding fixed the aggregate factor supply, changes in land-use policy would induce changes in regional allocations that would mitigate the adverse effects of tighter policies. Both models feature perfectly mobile capital across regions, but they differ in the mechanism used to allocate labor across space.

Perhaps the most unusual feature of the HOP model is the representative family construct used to determine aggregate labor supply and allocate workers across space. This essentially amounts to equating the marginal utility of workers across regions instead of the level. In many other papers studying the spatial allocations of people, the level of utility is usually equated across space (subject to moving costs) since it is the individual deciding where to locate, not the representative head of the family. The representative household can internalize externalities, generating a higher level of aggregate utility than the competitive equilibrium with individual decision makers, but it is not clear if the elasticity of utility with respect to land-use regulations is higher or lower. Perhaps the representative household is less sensitive to changes in land-use regulation because its household members are already partially insured against regional differences, whereas individual decision makers would be more elastic since they have no insurance against local shocks. The importance of this difference is an interesting area for future investigation. A first step could be to compare the level of utility of workers across regions using a region-specific utility function at the allocations of the baseline model as a quick and dirty way to see if the level of utility across space is similar as would be in the more typical spatial equilibrium setup.

Finally, one other difference could simply be the imputed magnitude of the land-use distortions, in which the models would predict similar predictions for output in response to a given change in land-use policies. HOP filter the change in land-use regulations and perform counterfactuals by changing a structural model parameter that affects the supply of a fixed factor of production. HM freeze house prices at 1964 levels. While α is closely related to house prices, it is not obvious how to compare the magnitude of the two different policies that define the counterfactuals, making it complicated to compare the elasticities of output to the respective policy changes across models.

4. Computing counterfactuals: reverse causality and heterogeneity

The paper assumes that the imputed tightening of land-use regulations is bad. I believe the data permits an alternative interpretation in which an increase in land-use regulations could increase output in a region. This primarily raises concerns about how to compute counterfactuals. That is, when changing α what is held fixed and what else changes.

Please cite this article as: C. Tonetti, Comment on "Tarnishing the golden and empire states: Land-use regulations and the U.S. economic slowdown," by Herkenhoff, Ohanian, and Prescott, Journal of Monetary Economics (2017), https://doi.org/10.1016/j.jmoneco.2017.11.002

The first version of this concern is partially addressed in Section 8.5. What if land-use regulation generates low-density housing that is valued as an amenity? If this were the case, then amenities should be changed in the counterfactual when changing α. To quantify this relationship, HOP use the correlation between state-level amenities and land-use regulations, which suggests that stricter land-use regulation is positively correlated with better amenities, to worsen amenities when counterfactually raising α. In practice, accounting for this relationship in this fashion does not have a large impact on the counterfactual. It is still possible that the relationship between amenities and land-use regulation varies across regions or time and may be a nonlinear function of population density. For example, when the population density of a region is low, it is likely that amenities could remain constant as the population increases, but when a region is already very crowded, any increase in population could severely degrade amenities due to congestion externalities. Thus, the endogenous relationship between amenities and land-use regulation seems an exciting open area for further investigation.

Another alternative interpretation of the historical trends revolves around the mobility of productivity. To illustrate this version of the concern, I will add one extra feature to the model: heterogeneous workers. If the true data generating process had two types of workers, with either high or low productivity, then locations with high measured TFP in the HOP model could be locations with large concentrations of high productivity workers. Further, if low-density housing is a normal good, then high productivity workers may spend part of their higher income to purchase low-density housing in the form of stringent land-use regulations. The purchase of low-density housing may be related to land-use regulation, as opposed to being reflected purely in private market transactions, due to externalities in land use. Thus, California may be a high (measured) TFP state *because* it has strict land-use regulations that attract high productivity workers. This potential direction of causation substantially changes counterfactual predictions: If California returned to the lax land-use policies of its past, it might loose the most productive workers, lowering measured TFP and output.

Even more broadly, any interpretation of the HOP results depends crucially on whether productivity is specific to a location or if it is mostly embedded in mobile factors of production (like workers). The position in the HOP paper is that productivity is location specific, so land-use regulations prevent homogeneous workers from migrating to high productivity locations. Analogous to the cross-country development accounting literature, perhaps the location specific productivity measures capture differences in state laws and institutions, as well as productive ideas that are difficult to move due to excludability and adjustment costs (e.g., a high quality idea is patented and used by a firm that paid a large fixed cost to build a headquarters in a particular region). From the polar opposite perspective, if instead productivity is completely embedded in mobile factors of production that are relatively cheap to move across space, the spatial correlation of TFP, amenities, and house prices may reflect very different forces than those captured by the HOP model. A model with heterogeneous workers may impose certain restrictions on the relationships between these unobservables, which may point towards or rule out such human-capital related stories of productivity mobility. Quantitatively examining the plausibility of these alternative interpretations seems like a fruitful area for future research.

5. Concluding thoughts

This paper takes an important step forward, using theory to fill in the gaps in historical data to provide a formal quantitative analysis of land-use regulations at the scale of the aggregate U.S. economy over decades. The analysis suggests reforming land-use regulations may be a powerful tool for raising output and welfare. There are a few unique features to the HOP analysis not shared by other papers that examine similar questions and it would be interesting in future research to see which these features matter most in generating differing elasticities of output to policy changes, with a focus on exploring direct evidence on these differences. Lastly, there is a wide-open question on how to perform counterfactuals that hinges on understanding the determinants of local measured productivity, namely whether productivity is determined predominantly by mobile or immobile factors.

References

Glaeser, E.L., Gyourko, J., Saks, R., 2005a. Why have housing prices gone up? Am. Econ. Rev. P&P 95 (2), 329–333.
Glaeser, E.L., Gyourko, J., Saks, R., 2005b. Why is manhattan so expensive? J. Law Econ. 48 (2), 331–369.
Glaeser, E.L., Ward, B.A., 2009. The causes and consequences of land use regulation: Evidence from greater boston. J. Urban Econ. 65 (3), 265–278.
Herkenhoff, K.F., Ohanian, L.E., Prescott, E.C., 2017. Tarnishing the golden and empire states: Land-use regulations and the U.S. economic slowdown. J. Monet. Econ.
Hsieh, C.-T., Moretti, E., 2017. Housing constraints and spatial misallocation. Working paper.
Kiyotaki, N., Michaelides, A., Nikolov, K., 2011. Winners and losers in housing markets. J. Money, Credit Banking 43 (2–3), 255–296.

Please cite this article as: C. Tonetti, Comment on "Tarnishing the golden and empire states: Land-use regulations and the U.S. economic slowdown," by Herkenhoff, Ohanian, and Prescott, Journal of Monetary Economics (2017), https://doi.org/10.1016/j.jmoneco.2017.11.002