

The Geography of Worker and Job Flows: A Rejoinder to Bilal (2022)

Moritz Kuhn* Iouri Manovskii† Xincheng Qiu‡

September 29, 2022

1 Literature Background

There are large and persistent differences in unemployment rates across local labor markets. The standard approach to modeling unemployment involves search and matching. [Lkhagvasuren \(2012\)](#) and [Kline and Moretti \(2013\)](#), among others, proposed spatial versions of the Diamond-Mortensen-Pissarides model with exogenous separations to rationalize spatial differences in unemployment. By construction, these models attribute the entire spatial unemployment differences to differences in job-finding rates across locations. This approach was criticized by [Bilal \(2021\)](#) who documented that spatial differences in job-separation rates are more important than differences in job-finding rates in accounting for spatial unemployment differences. To account for this empirical observation, he proposed a model of spatial employer sorting. He demonstrated that this model accounts for spatial differences in unemployment and yields a decomposition of the importance of job-finding and separation rates in line with the data.

The job-finding and separation rates—the focus of the existing literature—refer exclusively to *worker* flows between employment and unemployment, hence we refer to them as facts on the *worker side* of the labor market. Perhaps surprisingly, the existing literature has not considered the facts on the *employer side* of the labor market, in particular, the spatial differences in vacancy creation and vacancy filling. Local unemployment is an equilibrium outcome shaped by actions of both workers and firms. Spatial differences in vacancy creation and vacancy filling are obviously central to determining local unemployment differences, even in models that are exclusively motivated by capturing the properties of worker flows. Yet, prior to [Kuhn, Manovskii, and Qiu \(2021\)](#), the facts on the geography of vacancy creation and vacancy filling

*University of Bonn, Department of Economics. Email: mokuhn@uni-bonn.de.

†University of Pennsylvania, Department of Economics. Email: manovski@econ.upenn.edu.

‡University of Pennsylvania, Department of Economics. Email: qiux@sas.upenn.edu.

have not been documented. Consequently, an analysis of what theory is consistent with the full set of empirical facts was also missing.

2 Contribution of Kuhn, Manovskii, and Qiu (2021)

Kuhn, Manovskii, and Qiu (2021) aimed to fill these gaps in the literature and offered three main contributions:

1. They document new facts on the geography of vacancy creation and vacancy filling. They find that there are more vacancies per unemployed worker in high-productivity/low-unemployment locations and that it is more difficult and takes longer for employers to fill vacancies in high-productivity/low-unemployment locations.
2. They combine the classic Rosen (1979)-Roback (1982) spatial equilibrium model with productivity differences across locations and a standard DMP model with endogenous separations (Mortensen and Pissarides, 1994) within each location. They find that the resulting combination of the two separately widely studied textbook models quantitatively matches all the facts on the geography of job-finding, separation, vacancy-posting, and vacancy-filling rates.
3. They show that the empirical properties of on-the-job search help reconcile the different relative importance of separation and job finding in the cross-section of local labor markets and over the business cycle.

3 Modeling Approach of Bilal (2021)

Bilal (2021) explores a different model in which new jobs are born unattached to a specific location, next their productivity is revealed, and finally they pick one among many heterogeneous locations. Thus, this is a model of assortative matching between jobs and locations featuring two main dimensions of sorting.

First, jobs that are revealed to be more productive have a higher opportunity cost of being vacant. Thus, the timing assumption in the model gives rise to sorting on the probability of filling a vacancy: firms that have high expected productivity at birth will move to and post vacancies in a location where vacancies are filled faster. Introducing this mechanism is the main theoretical contribution of Bilal (2021) and the purported quantitative importance of this mechanism is argued central to the results in the paper, the efficiency properties of the model, and its policy implications. However, despite the fact that at the heart of this sorting

mechanism is the speed of filling a vacancy, [Bilal \(2021\)](#) does not contrast the core implication of this mechanism with the data. This mechanism implies that high-productivity locations (which is endogenous in the model and in part reflects sorting of jobs) have high vacancy-filling rates. The empirical findings in [Kuhn, Manovskii, and Qiu \(2021\)](#) show that these implications are counterfactual. In the data, higher productivity locations have lower vacancy-filling rates.

The second sorting mechanism in [Bilal \(2021\)](#)'s model is more conventional in the literature and is driven by the complementarity between location-specific exogenous productivity and idiosyncratic productivity of jobs choosing to locate there.¹ This induces positive assortative matching where more productive jobs sort into more productive locations, so that the *total* productivity differences across locations reflect the interaction between location-specific productivity and productivity of jobs sorting into those locations. [Kuhn, Manovskii, and Qiu \(2021\)](#) take total productivity differences across locations as given and show that their spatial equilibrium model with DMP local labor markets, endogenous separations, and free-entry into each location matches all the data, including vacancy-filling rates. Note that it is immaterial to the findings in [Kuhn, Manovskii, and Qiu \(2021\)](#) whether the productivity differences are the property of locations, or of the firms operating in those locations, or some combination of the two. In particular, it is irrelevant for their positive analysis whether total location productivity differences are exogenous or are induced endogenously by, e.g., firm sorting based on the complementarity between their and location's productivities. As discussed below, the equilibrium determination in the sorting model in [Bilal \(2021\)](#) is fundamentally different. Yet, the findings in [Kuhn, Manovskii, and Qiu \(2021\)](#) suggest that if sorting were based only or mostly on the complementarity between location and job productivity, this sorting mechanism might have been able to generate vacancy filling rates that decline in total location productivity.

Thus, [Bilal \(2021\)](#) combines two sorting mechanisms, one of which induces a positive co-movement between total location productivity and job-filling rates, and the other of which might induce a negative co-movement. With two mechanisms in play, the net result is theoretically ambiguous and needs to be confronted with the data. Unfortunately, [Bilal \(2021\)](#) does not report the fit of the estimated model to empirical vacancy-filling rates. However, as we explain below, the additional results that have now been reported in [Bilal \(2022\)](#) and in the appendix to the September 2022 revision of [Bilal \(2021\)](#) reveal that in his estimated model vacancy-filling rates increase in location's productivity, in direct contrast to the robust empirical properties documented in [Kuhn, Manovskii, and Qiu \(2021\)](#).

¹[Duranton and Puga \(2004\)](#) review theoretical microfoundations of local productivity differences; [Behrens, Duranton, and Robert-Nicoud \(2014\)](#) and [Gaubert \(2018\)](#) decompose productivity differences across space into entrepreneurs/firm sorting and agglomeration economies.

4 A Comparison of Modeling Approaches in Kuhn, Manovskii, and Qiu (2021) and Bilal (2021)

Recently, we became aware of a note² (Bilal, 2022) claiming that the model in Bilal (2021) nests the one in Kuhn, Manovskii, and Qiu (2021). If this were the case, given a near-perfect match of the latter model to the data, a flexible estimation procedure applied to the model in Bilal (2021) could have shut down endogenous sorting on vacancy-filling rates and matched the data through local productivity differences only (as the results in Kuhn, Manovskii, and Qiu (2021) indicate). Why doesn't it do so? Because the two models are actually not nested.

Bilal (2022) argues that just as Kuhn, Manovskii, and Qiu (2021), he proposes a DMP model with endogenous separations in each location nested into a Rosen-Roback spatial setting. In particular, he claims that in Bilal (2022), “within every location, workers and firms meet in a frictional labor market as in the DMP model.” While the DMP model does feature a frictional labor market, so do many other but very different models. What is central to the DMP model is its equilibrium determination of vacancy posting, typically through the assumption of free entry, which implies that vacant jobs are created until the value of creating a vacancy is driven down to zero. Kuhn, Manovskii, and Qiu (2021) maintain this equilibrating force *for every local labor market* to endogenously determine the equilibrium number of vacancies in each location, and show that this model generates properties of vacancy posting and filling across locations consistent with the data. In contrast, in Bilal (2021) there is no free entry into each location. Instead, his model is one of assortative matching where the two sorting mechanisms described above induce a single-crossing condition that drives sorting of vacant jobs of different productivity to locations and wages sustain this sorting. In other words, conditional on sorting, the number of jobs in each location is determined by the exogenous distribution from which vacant jobs draw their productivity. Consequently, despite appearing superficially similar to a DMP model in each location—simply because of featuring vacancies, unemployment, and job search—the model in Bilal (2021) is fundamentally different from the equilibrium determination of vacancies, unemployment, and wages in a DMP model, which is in fact the essence of the DMP model.³

²The note is dated September 1, 2022. It was retrieved on September 11, 2022.

³A simple analogy might be helpful here. Clearly, the DMP model and Shimer and Smith (2000) model of assortative matching between heterogeneous workers and jobs feature a frictional labor market but have very different equilibrium properties. The total mass of jobs in Shimer and Smith (2000) is exogenously fixed and the distribution of jobs of different productivity is fixed by the distributional assumption. Hagedorn, Law, and Manovskii (2017) add the ex-ante entry stage to the Shimer and Smith (2000) model where firms pay a cost of entry prior to learning the realization of their productivity draw. This endogenizes the total mass of firms but the distribution of firm productivity is still exogenously fixed by the distributional assumption. The lack of free entry type-by-type is essential to sustain sorting in the Shimer and Smith (2000) model, which will collapse if the model allowed for free entry of jobs conditional on their productivity type. Kuhn, Manovskii, and Qiu (2021) follow the DMP framework and allow for free entry of jobs into each location, and this determines the

Of course, a flexible enough estimation procedure applied to Bilal (2021)’s model could still have shut down sorting by, e.g., collapsing the distribution of idiosyncratic job productivity to a point, so that all jobs are homogeneous.⁴ This is precluded by a hardwired assumption ostensibly used by Bilal (2021) to “identify” the importance of spatial job sorting. Recall that total location productivity in Bilal (2021) is determined by an exogenous location-specific component and the productivity component specific to jobs choosing that location. Bilal (2021) imposes that workers’ flow utility of unemployment scales one-to-one with the exogenous location-specific productivity but it is not affected at all by the endogenous job component of productivity induced by job sorting.⁵ This assumption allows Bilal (2021) to “identify” job-quality heterogeneity through local separation rates (see his Equation (23) that characterizes a one-to-one relationship between the local job type and local separation rate). In the data, separation rates increase strongly with local unemployment. If flow utility of unemployment is proportional to location productivity, differences in local productivity cannot induce this pattern.⁶ Kuhn, Manovskii, and Qiu (2021) show that if workers’ flow utility of unemployment is not increasing in total productivity of locations, their model quantitatively matches the empirical pattern of the separation rate declining in local productivity. This can only be accomplished in Bilal (2021)—given the identifying assumption—by assigning a dominant role to endogenous job sorting in determining differences in total productivity across locations. In fact, Bilal (2021) directly pins down the importance of sorting by requiring that the model matches the differences in separation rates across locations and concludes that sorting is very important. However, this conclusion is circular: Assume that only firm sorting can induce the large observed differences in separation rates across locations and then conclude that sorting is important because separation rate differences across locations are large. Obviously, the model in Kuhn, Manovskii, and Qiu (2021) matches the data without featuring the sorting mechanism but the assumptions in Bilal (2021) rule this alternative possibility out by precluding nesting the model in Kuhn, Manovskii, and Qiu (2021). The assumption of a very important role of spatial job sorting in Bilal (2021) may not have been particularly problematic if sorting was based on the productivity complementarity only. However, the two sorting mechanisms are bundled in the model and assuming a large role of sorting on the probability to fill a vacancy is concerning as it yields

equilibrium properties of local vacancies. In contrast, the model in Bilal (2021) is analogous to Hagedorn, Law, and Manovskii (2017), although sorting is frictionless and is not between jobs and workers but between jobs and locations.

⁴We are not suggesting that there is no firm heterogeneity in the data or that firms are not sorting across locations based on productivity complementarities. Unfortunately, eliminating the role of sorting on the probability of filling a vacancy without eliminating sorting on productivity at the same time cannot be achieved in the framework in Bilal (2021) through a change in parameter values. Doing so requires changes to the model.

⁵Our intention is not to debate the assumption per se, but to highlight its implications in Bilal (2021).

⁶In fact, this will induce the opposite pattern of separation rates increasing in productivity and decreasing in local unemployment: high productivity locations should still feature higher job-finding rates, making workers more selective in the matches they accept, leading to an increasing separation threshold.

counterfactual implications for the co-movement of job-filling rates with productivity.

Note that the model [Bilal \(2021\)](#) actually takes to the data has many more components than discussed in the theory part of the paper, including amenities, endogenous recruiting intensity, permanent worker heterogeneity, human capital accumulation, scarring effects of unemployment, etc. These additional components provide further degrees of freedom in matching the data, but the role of the different mechanisms becomes much less transparent. Perhaps some or the combination of these additional mechanisms could help overcome the counterfactual implications for vacancy-filling rates induced by firm sorting on vacancy-filling rates, but unfortunately, this is neither discussed nor is there any evidence of this presented in the paper.

Finally, [Kuhn, Manovskii, and Qiu \(2021\)](#) offer a solution to the puzzle [Bilal \(2021\)](#)'s empirical finding poses, but the paper does not answer—differences in job-finding rates account for the majority of the unemployment fluctuations over the business cycle, whereas differences of separation rates account for the majority of the unemployment differences in the cross-section of local labor markets. It poses a theoretical challenge to generate both patterns in a unified framework. In the data, the job-to-job rate is strongly pro-cyclical over time (business cycle), whereas it does not systematically vary with the unemployment rate across space. [Bilal \(2021\)](#) concludes from the spatial fact that on-the-job search can be abstracted from because it does not vary systematically with local unemployment. [Kuhn, Manovskii, and Qiu \(2021\)](#) take an opposite route, model on-the-job search, and find that this empirical property of on-the-job search is actually key to reconciling the flow decomposition across space vs. over the business cycle.

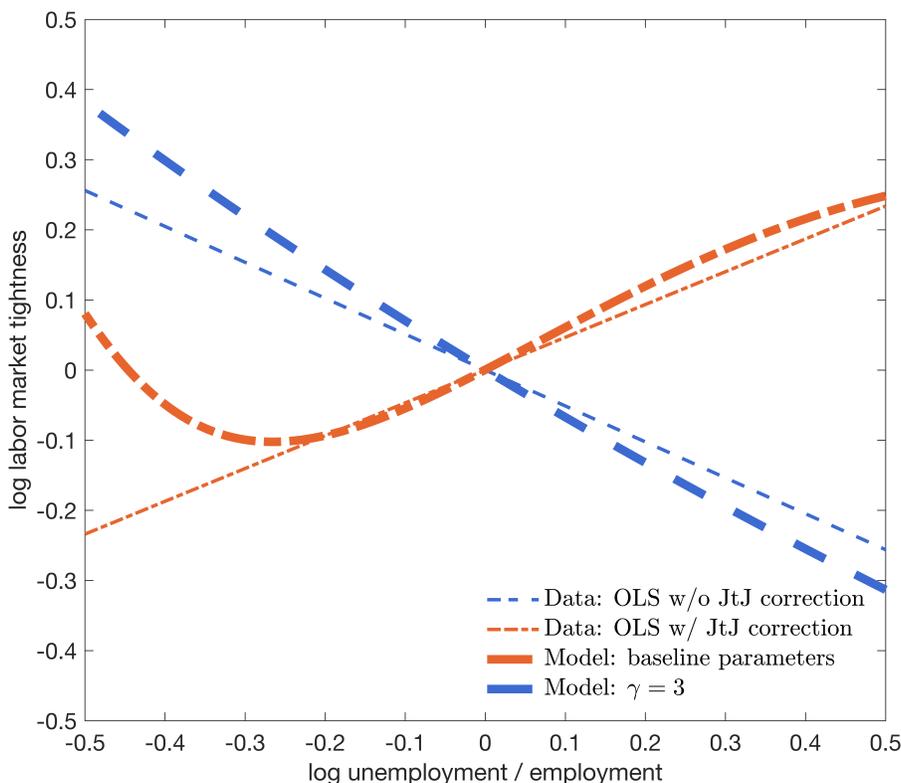
To summarize, [Kuhn, Manovskii, and Qiu \(2021\)](#) document key missing local labor market facts and present a model that is quantitatively consistent with all documented aspects of the data. A sorting model in [Bilal \(2021\)](#) is successful in accounting for the facts on the worker side of the market but it remains to be seen whether it matches the other relevant aspects of the data, including vacancy-filling rates. Although [Bilal \(2021\)](#) is silent on this, in the next section we will argue that the new information revealed in [Bilal \(2022\)](#) implies that it does not.

5 Is the Estimated Model in Bilal (2021) Consistent with the Data on Vacancies?

The motivation for [Bilal \(2021\)](#) was to account for the geography of worker flows between employment and unemployment. As was standard in the literature, the paper did not consider data on vacancy flows. Recently, presumably in response to [Kuhn, Manovskii, and Qiu \(2021\)](#)—which is not acknowledged in [Bilal \(2021\)](#) despite being available and widely presented for over

Figure 1: Figure 1 Reproduced from Bilal (2022)

Figure 1: Labor market tightness against unemployment-to-employment ratios in the model in Bilal (2021) and in the data.



Note: Figure 1 plots the log of labor market tightness at baseline parameters and when setting $\gamma = 3$ in the model in Bilal (2021). It also plots regression lines from regressing labor market tightness onto the log unemployment-to-employment rate in French data, consistently with Figure 6 in Bilal (2021). Figure 1 reports the regression lines with and without the job-to-job correction in the data.

two and a half years, the paper added some discussion on the patterns of labor-market tightness across local labor markets in France but not the fit of the model to these patterns. The model fit was first shown in Bilal (2022) and added as an appendix figure in September 2022 revision of Bilal (2021). This plot is reproduced in Figure 1.

The decision to plot tightness rather than the model prediction for vacancy-filling rates is somewhat perplexing because it is the vacancy-filling rates that drive the novel sorting mechanism in the model in Bilal (2021) and the concern that precisely because of this sorting mechanism the model will deliver the co-movement between vacancy-filling rates and local productivity (or unemployment) of the opposite sign relative to the data. Moreover, the measurement of vacancy-filling rates in the data is direct and does not require any adjustments for on-the-job search that Bilal (2021) suggests might cloud the interpretation of the measured labor market tightness. Nevertheless, we can infer the relationship between vacancy-filling rates

and local unemployment (and productivity) in the estimated model in Bilal (2021) from the newly revealed information on labor market tightness in that model.

The prediction for the co-movement between labor market tightness and local unemployment in the estimated model in Bilal (2021) is reported as a thick orange line in Figure 1. The relationship is generally upward sloping suggesting that there are more vacant jobs per unemployed worker in labor markets with higher unemployment. This pattern is claimed to be a success of the model as it aligns reasonably well with the particular empirical measure of labor market tightness adjusted for on-the-job search plotted as an orange thin line.⁷

To understand the implication of this figure for vacancy-filling rates, note that in Bilal (2021) the vacancy-filling rate π^{ve} is comprised of two components: the contact rate for vacant job q and the probability that the contact turns into a match a . All variables could vary by location j so that $\pi_j^{ve} = q_j a_j$. Similarly, the job-finding rate π^{ue} is also composed of two components: the contact rate for an unemployed worker p and the probability that the contact turns into a match a , such that $\pi_j^{ue} = p_j a_j$. Note that the probability that a contact turns into a match cancels out when we take the ratio of the vacancy-filling rate and the job-finding rate,

$$\frac{\pi_j^{ve}}{\pi_j^{ue}} = \frac{q_j}{p_j} = \frac{1}{\theta_j},$$

where $\theta_j := v_j/u_j$ and the second equality holds for any constant-returns-to-scale matching function. Given the prediction that θ is increasing in local unemployment in the estimated model in Bilal (2021) as reported in Figure 1 and a nearly-flat job-finding rate as claimed in Bilal (2021), this immediately implies that the vacancy-filling rate π_j^{ve} is decreasing in local unemployment. This model prediction is inconsistent with the empirical evidence in Kuhn, Manovskii, and Qiu (2021). In fact, given that the job-finding rate is slightly decreasing in local unemployment in Bilal (2021), the vacancy-filling rate has to decrease even more strongly according to the derived relationship.

⁷The discussion of how to adjust labor market tightness for on-the-job search is rather superfluous, because the truly relevant measure of the vacancy-filling rates is not affected by this adjustment (it measures the rate at which the vacancy is filled regardless of whether the new hire came from unemployment or from another job). Nevertheless, we note that Kuhn, Manovskii, and Qiu (2021) have presented robust evidence on both the more conventional measure of labor market tightness defined as the v/u ratio as well as OJS-adjusted labor market tightness defined as $v/(u + \phi e)$, where ϕ represents the relative search intensity of the employed workers. For the adjustment, Kuhn, Manovskii, and Qiu (2021) leverage direct information from data on new hires indicating whether a new hire was previously unemployed or employed in another job (see their Appendix I.1.5 for details). The adjustment in Bilal (2021) is unorthodox (he sets $\phi = 0.92$). The implied increasing pattern in local unemployment is also surprising: OJS adjusted tightness lies between v/u (when $\phi \rightarrow 0$) and v (when $\phi \rightarrow 1$). Both v/u and v are downward sloping in local unemployment (the latter relationship is known as the Beveridge curve) in the data in Kuhn, Manovskii, and Qiu (2021) and in the French data reported in Fournier (2021), albeit for Paris region only. This makes the finding that something between v/u and v is upward sloping unexpected.

The logic could be easily extended even if one introduces recruiting effort, denoted by s , so that

$$\frac{\pi_j^{ve}}{\pi_j^{ue}} = \frac{s_j q_j}{p_j} = \frac{s_j}{\hat{\theta}_j} = \frac{1}{\theta_j},$$

where $\theta_j := v_j/u_j$, and $\hat{\theta}_j := s_j v_j/u_j$ is defined as the effective tightness that adjusts for recruiting intensity s_j . It is not clear whether Bilal (2022) is plotting θ or $\hat{\theta}$ as labor market tightness in his model prediction in Figure 1. If what he plots is $\theta_j := v_j/u_j$, then the logic of the previous paragraph applies directly, implying that the vacancy-filling rate π_j^{ve} is decreasing in local unemployment in the model in Bilal (2021). We suspect that the figure might actually show $\hat{\theta}_j := s_j v_j/u_j$, i.e., effective tightness adjusted for recruiting intensity s_j .⁸ Nevertheless, even if the figure plots $\hat{\theta}_j$, it would imply that the vacancy-filling rate π_j^{ve} is decreasing in local unemployment even more, as s_j is decreasing in unemployment (see Equation (71) in the online appendix of September 2022 revision of Bilal (2021) that recruiting effort is monotone in the value of a filled job).

Thus, from the partial information Bilal (2022) reveals about his model performance on tightness, it seems pretty clear that even his full model under the estimated parameter is counterfactual. In his model, vacancies are filled faster the lower unemployment in a location is. In the data, vacancies are filled faster in higher unemployment locations.⁹

6 Conclusion

It seems absolutely clear that Bilal (2021) was motivated by trying to get the spatial differences in worker-flow rates right. His work was successful in documenting the relevant facts on the relative importance of job-finding and separation rates across local labor markets. He also proposed a successful theoretical model matching the properties of these flows. This was a very valuable and influential contribution to the literature.

⁸It is unclear how to compare effective tightness in the model to the tightness measured in the data. In the data, vacancies are measured as the number of unique job openings looking to hire. Even in the vacancy data coming from, say, online adds, a meticulous effort is taken to remove multiple postings on different websites that refer to the same vacant job. Measuring the effective vacancies is akin to measuring the number of websites on which each job is advertised rather than the actual number of vacant jobs.

⁹This makes the accusation in Bilal (2022) that Kuhn, Manovskii, and Qiu (2021)’s discussion of Bilal (2021)’s novel mechanism focuses on “a very special case of the model Bilal (2021)” that “assumes away any exogenous local productivity differences” rather jarring. It turns out that the concerns expressed Kuhn, Manovskii, and Qiu (2021) about the counterfactual implications of the mechanism of sorting on vacancy-filling rates are quite profound as they outweigh all other mechanisms at the estimated parameter values in Bilal (2021) and are featured by the full model on which the quantitative analysis is based and which is used to study the effects of policies. It may or may not be the case that in Bilal (2021), “depending on parameter values, the model can generate any correlation between vacancies, labor market tightness, vacancy filling rates.” The relevant question is what happens at the estimated parameter values. We now know the answer to this question.

The objective of [Kuhn, Manovskii, and Qiu \(2021\)](#) was different. They recognized that the literature had not yet documented the facts on the employer side of the labor market, in particular the spatial differences in vacancy posting and vacancy filling. Having documented these facts—an important and nontrivial exercise in and of itself—they proceeded to develop the first quantitative model that is simultaneously consistent with the facts on both worker and employer sides of local labor markets.

Thus, [Bilal \(2021\)](#) is not a competitor but an important stepping stone for [Kuhn, Manovskii, and Qiu \(2021\)](#). It is sad to see an attempt in [Bilal \(2022\)](#) to rewrite history and argue that [Bilal \(2021\)](#) has also aimed to match the facts on both the worker and employer sides of the market. Not only [Bilal \(2021\)](#) never had that objective, it had not even reported the differences in vacancy-filling rates in the data or in his model. If [Bilal \(2022\)](#) correctly captured the intentions of [Bilal \(2021\)](#), would not [Bilal \(2021\)](#) have reported the vacancy-filling rates in the quantitative model illustrating the empirical support for the core mechanism in his theory, laying to rest the concerns expressed in [Kuhn, Manovskii, and Qiu \(2021\)](#), and positioning his model as a potential alternative to the model in [Kuhn, Manovskii, and Qiu \(2021\)](#)? Only because [Bilal \(2021, 2022\)](#) never did so, did we have to infer the patterns in his estimated model ourselves. As our analysis revealed, the quantitative predictions of his estimated model are in direct conflict with the patterns of vacancy-filling rates as documented in [Kuhn, Manovskii, and Qiu \(2021\)](#). While this might be construed as a critique of [Bilal \(2021\)](#) as a model accounting for both worker and employer sides of the labor market, if [Bilal \(2021\)](#) is instead assessed as a model of the worker side of the market only—as we think it should be—one will recognize that it is a very valuable contribution with important and original insights.

References

- Behrens, Kristian, Gilles Duranton, and Frédéric Robert-Nicoud. 2014. “Productive Cities: Sorting, Selection, and Agglomeration.” *Journal of Political Economy* 122 (3):507–553.
- Bilal, Adrien. 2021. “The Geography of Unemployment.” Tech. Rep. 29269, National Bureau of Economic Research.
- . 2022. “Theories of Job Loss, Job Finding and Vacancies in the Cross-Section of Locations: A Response to Kuhn et al. (2021).” Tech. rep., Harvard University.
- Duranton, Gilles and Diego Puga. 2004. “Microfoundations of Urban Agglomeration Economies.” In *Handbook of Urban and Regional Economics, Vol. 4*, edited by J. V. Henderson and J.-F. Thisse. Amsterdam: Elsevier-North Holland, 2063–2117.
- Fournier, Juliette. 2021. “Spatial Mismatch.” *Unpublished manuscript* .
- Gaubert, Cecile. 2018. “Firm Sorting and Agglomeration.” *American Economic Review* 108 (11):3117–53.
- Hagedorn, Marcus, Tzuo Hann Law, and Iourii Manovskii. 2017. “Identifying Equilibrium Models of Labor Market Sorting.” *Econometrica* 85 (1):29–65.
- Kline, Patrick and Enrico Moretti. 2013. “Place Based Policies with Unemployment.” *American Economic Review, Papers and Proceedings* 103 (3):238–243.
- Kuhn, Moritz, Iourii Manovskii, and Xincheng Qiu. 2021. “The Geography of Job Creation and Job Destruction.” NBER Working Papers 29399, National Bureau of Economic Research.
- Lkhagvasuren, Damba. 2012. “Big Locational Unemployment Differences Despite High Labor Mobility.” *Journal of Monetary Economics* 59 (8):798–814.
- Mortensen, Dale T. and Christopher Pissarides. 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies* 61 (3):397–415.
- Roback, Jennifer. 1982. “Wages, Rents, and the Quality of Life.” *Journal of Political Economy* 90 (6):1257–1278.
- Rosen, Sherwin. 1979. “Wage-Based Indexes of Urban Quality of Life.” *Current Issues in Urban Economics* :74–104.
- Shimer, Robert and Lones Smith. 2000. “Assortative Matching and Search.” *Econometrica* 68 (2):343–369.