Wage Variations and Commuting Distance

El-Mehdi Aboulkacem [1]

Clément Nedoncelle *[2]

[1] UGE – AME - SPLOTT, 59666, Villeneuve d'Ascq Cedex, France
[2] Université Paris-Saclay, INRAE, AgroParisTech, Paris-Saclay Applied Economics, 91120, Palaiseau, France.

* Corresponding Author. Email: clement.nedoncelle@inrae.fr

Abstract:

We estimate the causal impact of wage variations on commuting distance of workers. We test whether higher wages across years lead workers to live further away from their working place. We use employer-employee data for the French Ile-de-France region (surrounding Paris), from 2003 to 2008, and we deal with the endogenous relation between income and commuting using an IV strategy. We estimate that increases in wages coming from exogenous exposure to trade activities lead workers to increase their commuting distance and to settle closer to the city of Paris historical center. Our results cast novel insights upon the causal mechanisms from wage to spatial allocation of workers.

Keywords: Commuting; Wage; Distance; Trade**.**

JEL Classification: R2, R4, F6

Acknowledgments and Funding: We would like to thank the editor and the anonymous referees for the advice on this manuscript. We would also like to thank seminar participants in Lille, Paris and Toulouse, discussants in many workshops as well as conference participants and many colleagues for discussions. This work is supported by a public grant overseen by the French National Research Agency (ANR) as part of the « Investissements d'avenir » program (reference: ANR-10-EQPX-17 – Centre d'accès sécurisé aux données – CASD).

Data Availability Statement: The data that support the findings of this study are available from French Statistical Authorities. Restrictions apply to the availability of these data, which were used under license for this study. Data are available with the permission of CNIS (French Office for Statistical Secret).

Conflicts of Interest: The authors declare no conflict of interest.

Additional Material: This paper features an online appendix.

1. Introduction

In recent decades, the average home-workplace distance has increased enormously all over the world and not only in advanced economies. This rise in commuting distance resulted in several environmental issues such as air pollution, large use of fossil fuels or land artificialization among others. Overall, welfare is affected as inhabitants face pollution and traffic jams, and have to pay higher taxes to maintain the existing transport infrastructures and to build new ones. In this context, understanding the determinants of workers commuting distance is of great importance[1].

The existing theoretical literature, both in urban economics and in labor economics (reviewed hereafter), has focused on the relationship between income and commuting distance and provides ambiguous results, suggesting a mixed sorting of households in space along their income. The empirical literature on that specific topic also provides ambiguous and contradictory results (see Wheaton (1977); Timothy and Wheaton (2001); Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016)). While there is a robust correlation between wages and commuting distance that could result from many mechanisms, the sign and direction of the causality are not clearly established. Our paper tackles this issue.

The paper addresses the question of the causal impact of wage on the commuting distance at the individual level. In the present paper, we use wage variations across years within an individual worker to identify the direction of causality. In practical terms, do workers move further from their workplaces if their employer raises their wages? Our answer is: on average, yes. We use data that cover the universe of French exporting firms' employment in the *Ile-de-France* region (surrounding Paris) from 2003 to 2008. We use bilateral commuting distance and travel time data between all cities in the region. This dataset is matched with an exhaustive employer-employee administrative dataset providing information about workers' earnings, living and working cities. We estimate the worker-level commuting distance elasticity to wage.

The main empirical difficulty arises from the endogenous relationship between income and commuting. Previous attempts to control for the endogenous relationship between commuting and income exist. Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016) estimate a negative elasticity using an instrumental variables strategy in a two-stage least squares estimation to overcome endogeneity issues, together with sample restrictions. The main difficulty with this approach is to find suitable instruments for the income variable, that may affect the estimated elasticity[2]. Candidates are variables affecting income but not location (except through income). We

[1] These issues are far from new. For example, Rouwendal and Rietveld (1994) already stated in an article published in the early nineties that the ``increase in commuting distances is (…) a matter of concern for regional and national governments because of external effects in terms of congestion and environmental damage''. A quarter century later, these issues are still resonating.

[2] Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016) acknowledge that using alternative instruments for income affect the sign and the magnitude of the estimated elasticity; see their footnote 1.

use in the present paper a robust and convincing instrument for individual wage: we identify trade shocks that are exogenous to firms and workers decisions. In particular, we leverage trade shifts that are exogenous to firms and workers by computing, using disaggregated customs data, the world import demand that is addressed to each exporting firm, in the spirit of Autor et al. (2020) and Mayer, Melitz, and Ottaviano (2021) among others.

We take advantage of the plausible and documented observed randomness of these export shocks across workers as a source of variation for individual wages. In formal terms, we instrument individual wages with this measure of trade exposure to shocks in an IV-2SLS estimation of the distance-wage elasticity. We hence identify the wage variations, arising from trade exposure only, and not by other simultaneous shocks to firms or workers that may be related to location decisions. We provide results supporting the validity of this instrument. We thus settle in a recent literature using this shift-share type of instruments, in many contexts, including French firm-level data (for instance, Aghion et al. (2018) or Berman, Berthou, and Héricourt (2015)). Our strategy yet differentiates from this literature by instrumenting a worker-level variable with such a trade shifter.

We estimate that workers that experience higher wages, because of exports shifts, also live farther away from their workplace, ceteris paribus. Our fixed effect structure ensures that this effect occurs for a given worker across years: in our preferred estimates, a 10% increase in wages raises commuting distance by 2% on average in the short run. Then, using the sample of workers who effectively changed their home city, the elasticity goes up to 4% for a 10% increase in earnings. We interpret it as the long-run elasticity, as moving frictions are excluded from the estimation.

Finally, our results have implications regarding commuting patterns. As our strategy relies on a diff-in-diff estimation, our model allows us to study the location choice of the workers. By including information about *city* characteristics in our dataset, we can interpret the choice of workers in terms of economic characteristics of new home cities compared to former home cities. In particular, we find that workers that experience wage increases are moving closer to the center of Paris where historical amenities are abundant, where density is larger but where jobs are relatively scarce. Furthermore, the relative time cost of using private vehicles for home-workplace trips compared to public transportation is higher in the new home city than the former one. These results (i) support that amenities shape the spatial sorting of workers in cities, (ii) rule out the classic explanation of longer commutes for lower housing prices, and (iii) that in the particular case of Paris region, a higher income could result into a non-negligible modal shift from the private vehicle to public transport for home-workplace trips.

Our results are robust to different measures of wages and distance, to omitted variables, and to alternative instruments. In particular, we also checked the validity of the instruments, accounting for the potential non-random exposure of firms to foreign shocks (because of unobservables) and for potential minor violations of the exclusion restriction.

Our paper contributes to the literature in three dimensions. First, we provide a *causal* short-run estimate of the distance elasticity to wage, using an IV strategy. In particular, we leverage trade shifts uncorrelated to firms' and to workers' decisions to instrument for individual wage. While this type of instrument is extensively used in many contexts, we use it to estimate an under-studied elasticity in the urban economics literature. This contributes to the literature by providing the causal effect of wages variation on commuting distance, while existing results in the literature may either be flawed by reverse causality or omitted variable bias.

Second, our results provide alternative and additional estimates to the rare existing ones in the literature. In particular, Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016) estimate a negative causal elasticity, using IV and sample restrictions, in Denmark. With respect to their estimation, we use an alternative IV of individual wages and estimate a positive elasticity, around Paris. In this paper, we thus show that both different methodologies and different setups lead to mixed results.

Third, our data allow us to estimate two elasticities: a short-run elasticity (using the whole sample of workers) and a medium-/long-run elasticity (using the sample of workers that indeed changed residence, as in Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016), dropping workers which experience no residential move). These two elasticities differ as the short-run estimate accounts for moving costs and thus resistance to move: indeed, in the short-run, due to the presence of residential moving costs, even when experiencing a large change in household income, few workers will immediately change commuting distance by moving residence. However, conditioning the estimations on a residential move abstracts from these resistances and results can be interpreted as medium/long-run estimates. In practice, we either keep or drop residential stayers to obtain these two estimates. Our paper contributes to the literature as it provides both estimates using the same IV strategy.

The paper is structured as follows. The next section reviews the existing literature related to the research question. Section 3 presents the main datasets we use in the analysis and Section 4 describes the empirical strategy. Section 5 provides the main results. Section 6 details the robustness checks. The final section concludes.

2. Background

In this section, we review the existing literature on income and commuting distance.

2.1. Income and workers' sorting over space

In the urban economics literature, the first body of works builds upon the seminal urban land use model (Alonso (1964)). Fujita (1989) discusses many theoretical models in the basic simplest monocentric city framework. In this kind of model, workers are assumed to seek higher utility given a constant income by trading off proximity to employment against the attractiveness of residential sites. In equilibrium, higher-income households seek *further* residential locations to enjoy a larger and more pleasant living environment[3].

The empirical relevance of this sorting relies on the value of the elasticity of housing demand to income, as also highlighted by time-extended models. Old results suggest that this elasticity could be larger than one (Mills (1972)) -- meaning that as a household income increases, it is very likely to move to an area where housing is cheaper, namely further away from employment, in order to enjoy larger housing than what could be obtained if it stayed in the same area. Yet, this view is at odds with more recent evidence and empirical regularities. Glaeser, Kahn, and Rappaport (2008) or Rosenthal (2014) for instance provide evidence that the income elasticity of demand for land is far less than one, thus making this sorting implausible.

Another generation of works departed from this seminal urban land use model and identified additional determinants for the location of households with different incomes in cities. Among others, the literature has first emphasized the role of urban amenities as determinants of the location of rich households (Glaeser, Kolko, and Saiz (2001)). Brueckner, Thisse, and Zenou (1999) show that when the center has a strong amenity advantage relative to the suburbs' amenities, rich households are likely to live in central locations, instead of in the suburbs. Household sorting is thus determined by the location of amenities[4]. Recent literature (Duranton and Puga (2014); Koster, van Ommeren, and Rietveld (2014); Cuberes, Roberts, and Sechel (2019); Gaigné et al. (2022)) conclude that amenities do play an important role in the observed changes in demographic shifts (including neighborhood changes, gentrification)[5].

Then, beyond amenities, the age of the housing stock may generate household sorting within a metropolis as high-income households tend to locate in areas of the city where the housing stock is relatively young as evidenced in Brueckner and Rosenthal (2009). On top of these comes the effect of public transportation systems. Glaeser, Kahn, and Rappaport (2008) argue that the urbanization of poverty appears to be the result of better access to public transportation in central cities: public

[3] In the monocentric city models, land is a featureless plain. Only distance from city center/CBD differentiates locations.
[4] Note that the sorting over amenities may lead to the possibility of multiple equilibrium where the location of the rich is driven by accident, or history, but once established is relatively persistent, see for example Brueckner, Thisse, and Zenou (1999).
[5] For instance, Hoelzlein (2019); Guerrieri, Hartley, and Hurst (2013); McKinnish, Walsh, and Kirk White (2010).

transportation offers a time-intensive alternative that will be more appealing to those with low incomes. More recently, Tsivanidis (2018) studies Bogotá's response to the construction of the world's largest Bus Rapid Transit system. Among other consequences, the system generated a re-sorting of workers: the high-skilled workers moved into high-amenity, expensive neighborhoods while the low-skilled ones moved into poorer neighborhoods. Finally, de Bartolome and Ross (2003) show that house prices capture both the commuting advantage of cities and the fiscal difference between jurisdictions. As a result, income sorting across jurisdictions and across space are simultaneously at play, leading to a mixed sorting of households in space.

All these arguments explain why a strict income sorting is not generally observed in the data[6]: while initial evidence was in favor of expecting rich households out of the city center to enjoy cheaper land, recent evidence suggests that amenities, public transportation improvement, heterogeneous housing stock and heterogeneous local tax levels (and others) may create a mixed and dynamic sorting of households in space, with rich households potentially located within the city and close to downtown areas (Couture and Handbury (2017)).

2.2. Wages, Commutes, and Access to Jobs

Several mechanisms explored in the labor economics literature explain the relationship between commuting distance and wages. In particular, wages can be influenced by the length of commutes or more broadly by access to jobs. For that strand of literature, wages tend to increase with commuting distance for many reasons.

The first reason is for compensation and bargaining. Firms located far from workers' residences tend to compensate them for transport costs in order to attract them (Fujita, Thisse, and Zenou (1997); Timothy and Wheaton (2001); Mulalic, Ommeren, and Pilegaard (2014)). As a result, observed wage differences within cities could be offset by longer commutes and by compensation (Fu and Ross (2013)). On the other hand, workers won't accept jobs located far away from their homes unless the monetary gains they get are high enough to compensate for the travel monetary expenditure and the related loss of utility (Manning (2003)). Third, workers facing high commuting costs tend to bargain with their employers to get compensated (see Van Ommeren and Rietveld (2005); Mulalic, Ommeren, and Pilegaard (2014))[7].

[6] Note also that there is variation across cities and countries in this spatial sorting. Glaeser, Kahn, and Rappaport (2008) show that the households' income increases with distance from the CBD in the newer American cities, whereas it is not the case in the older ones. In France, Brueckner, Thisse, and Zenou (1999) document that richer households reside near the center while the poorest live in the suburbs in Paris and Lyon metropolitan areas, and that the opposite is the case in the average French city.

[7] Zenou (2009) supports that wages are set independent of the length of the commute in a standard monocentric urban model where employers have monopsony power, because house prices fully compensate for the length of the commute. In this context, at the equilibrium, workers with the same wages enjoy the same utility whatever their location in the city. However, in reality, cities are far from being monocentric and land use equilibrium is not always achievable. A perfect compensation of commuting monetary and non-monetary costs through house price is thus impossible

The second reason arises in efficiency-wage models. The positive spatial gradient is also documented in the literature in efficiency-wage models (Ross and Zenou (2008) for instance and Giménez-Nadal, Molina, and Velilla (2018)): long commutes affect leisure time at home and effort at work, affecting both wages and unemployment risks[8].

Another reason for the spatial gradient relates to discrimination in the labor market based on residential location. Among others, Gobillon, Selod, and Zenou (2007) and Jin and Paulsen (2018) show that some workers experience poor labor market outcomes because of their location (as they are disconnected from distant job opportunities for instance). A large part of the literature focuses on the adverse labor market outcomes of *minorities* because of their location, in line with the "spatial mismatch hypothesis" (Kain (1968), Arnott (1998)).

2.3. The role of broader economic shocks

Broad economic shocks simultaneously impact labor markets, wages, and commuting patterns. There is literature estimating the impact of local labor demand on labor markets (for a review, refer to Moretti (2011)), and recent emphasis is made on commuting. Monte, Redding, and Rossi-Hansberg (2018) show in a quantitative general equilibrium framework that both wages and commuting react to growth in local demand. Gyourko, Mayer, and Sinai (2013) show in the US that the increase in high incomes can explain the upward co-movement of incomes and house prices observed in "superstar cities." Couture et al. (2019) provide a spatial model of a city with heterogeneous agents and neighborhoods with endogenous amenity quality. A positive economic shock, that raises revenues, drives up house prices, and spurs the development of higher quality neighborhoods downtown, in line with evidence about increased commuting in Couture and Handbury (2017).

2.4. Empirical estimates and challenges

Most existing empirical results remain mainly silent regarding the causal impact of income on commuting and location decisions. The first set of existing results is directed towards the identification of the size and the sign of the effect. Early papers, such as Rouwendal and Rietveld (1994); Van Ommeren, Rietveld, and Nijkamp (1997); Van Ommeren, Rietveld, and Nijkamp (1999), study the relation between job change and residential change. On the other hand, Mulalic et al. (2014) and Timothy & Wheaton (2001) study how wages respond to the organization of cities and workers' locations.

There are many challenges to estimate the *causal* impact of wages on commuting. First, as wages are influenced by commutes, that represents a threat of reverse causality. Second, there are many

[8] In that strand of literature, wages are not the only labor market outcome to be affected by commuting.

omitted variables that would lead both to higher wages and longer commutes. Identifying causality is thus challenging, because of endogeneity concerns.

To the best of our knowledge, there are only a few papers that provide an estimate of the *causal* impact of wages (or income) on commuting patterns. Some papers, such as Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016), acknowledge both the existence of unobserved factors and reverse causality issues, and deal with this issue by using sample restrictions (only): they argue that reverse causation is less likely to be at play when workers move residence but do not change employer and thus estimate the elasticity on this sample of workers. As an alternative to using sample restrictions, a standard method for dealing with these endogeneity difficulties is IV estimation. The main problem with this approach is to find suitable instruments for the wage variable all the more that, as suggested by Manning (2003), the choice of the IV affects the estimated elasticity.

In this context, we argue (see Section 4) that world export shocks, translated into firm-level demand shocks, are a good candidate for a wage shifter. On the one hand, recent evidence supports that firm-level exports affect wages at the worker level, see e.g. Amiti and Davis (2012) or Carluccio, Fougère, and Gautier (2015). On the other hand, it is also hard to argue that export shocks are related to the commuting pattern at the individual level.

3. Data

We use datasets that cover the universe of French exporting firms' employment in the Ile-de-France, surrounding Paris, from 2003 to 2008. With 12 million inhabitants, Ile-de-France is the second European urban agglomeration (after Moscow) in terms of population and its largest employment area, with more than one million firms generating more than 6 million jobs. It is divided into 8 *départements* with a total of 1296 cities, spread across an area of 12,000 square kilometers. Paris consists of 20 districts (considered here as "cities" in the empirical analysis), is located at the center of the region, accounts for 1/3 of its jobs, and 1/3 of these jobs are located in the "historical" center of the city (1st to 9th district). As a result, only 1 job out of 9 in the region is located in the city of Paris "historical" center. Thus, the geographical framework we consider here is an urban space where attractive residential amenities are in general abundant in the area where jobs are scarce, namely in the historical center[9].

Worker wage and location data.

First, we use the *Déclaration Annuelle des Données Sociales* (DADS, « postes ») files. This matched administrative employer-employee dataset comes from the mandatory reports by firms about their

[9] Some areas of the region may be considered rural; However, Ile-de-France is as a whole an urban area similar in many respects to a very large metropolis.

workforce each year and is made available to researchers by the INSEE (*Institut National de la Statistique et des Etudes Economiques*). While it covers the full universe of the private sector, we only use the information on the individuals that worked once for a manufacturing firm. The unit of observation is a match between a firm (that can be identified) and an employee (that cannot be identified) over two years[10]. For each firm-worker observation, we have information about the individual gender, age, birth region, occupation (via 1 and 2-digit occupation code), annual gross and net earnings, the number of hours worked and the job status (full or part-time). Each worker-job observation also brings information about living and working places at the city level, using a 5-digit zip code.

Regarding working places, a firm may have many plants in the region (and potentially in France): in that case, we identify the 5-digit zip code of the city of the plant in which the worker effectively works at. As a result, some workers may stay in the same firm but in different plants over years. Our sample will account for this. Then, regarding the housing place, as it is defined at the city level, our data is unable to capture workers moving out of their homes but still living in the same city. In the empirical exercise, we will capture only residence changes across cities, neglecting within-city changes.

We apply some restrictions to the dataset. First, we restrict the information to workers both living and working in the IDF region. Second, we only keep full-time workers, to exclude obvious concerns regarding total earnings for workers with many part-time jobs. Third, since workers cannot be followed for more than two consecutive years, we restrict our sample to workers that stay within the same establishment (i.e. the same plant) across two consecutive years. We thus exclude job movers and restrict our sample to "job stayers". This reinforces prevention against reverse causality. Our results shall hence be interpreted with care: our estimation measures the impact of wage variation on the location decision of "job stayers". These workers represent more than 65% of the total earnings reported in the raw dataset.

Commuting Distance and Time data.

We merge this worker-level dataset with a unique dataset that provides information about the commuting distance and travel time across cities in the region. The dataset is provided by the *Société du Grand Paris* and informs about the effective commuting distance and travel time between cities in the Ile-de-France region. The data contains distance and travel time for the exhaustive bilateral set of cities in the region. Distance and travel time data, in kilometers and minutes respectively, are measured from city centroid to city centroid. This has many consequences: the

[10] Our results also neglect the possibility that firms may exit and close following trade shocks. We are however confident that producer entry and exit are of low magnitude in our sample. As we can identify all firms in our sample, we find that on average, our sample contains more than 85% of all firms that we present once in the raw dataset.

main one is that travel time and distance are equal to 0 for workers living and working in the same city[11].

Distance is our baseline measure of commuting. We obtain a set of time-invariant bilateral commuting distances and times that covers all potential travels made within the region. We merge this information for each worker, given time-varying home and working cities. We consider other measures, in both transport modes, in the sensitivity analysis.

Firm-level Trade Data.

We also make use of exhaustive firm-level trade data from the French customs to identify exporting firms and export shocks. This database reports the volume (in tons) and the value (in Euros) of exports for each CN8 product (European Union Combined Nomenclature at 8 digits) and destination, for each firm located in the French metropolitan territory. Some shipments are excluded from this data collection.[12] From this dataset, we only keep merchandise shipments, excluding agricultural and services exports. The raw dataset consists of 26,186,006 observations at the firm-year-destination-product level, which we aggregate into 7,110,894 observations at the firm-year-destination level and into 1,381,500 observations at the firm-year level.

Balance-Sheet Data.

We complete the picture using a balance-sheet dataset constructed from reports of French firms to the tax administration over the 2003-2008 period (*Bénéfices Réels Normaux, BRN*). This dataset contains information on the value-added, total sales, capital stock, debt structure, and other variables at the firm level. Importantly, this dataset is composed of all types of firms, including both small and large firms, since no threshold applies to the number of employees for reporting to the tax administration. This dataset allows us to control for firm characteristics in the empirical analysis.

Aggregate Product-level Trade Data.

We finally use the BACI dataset, providing country-level information about foreign import demand. BACI (Gaulier and Zignago (2010)) provides disaggregated data on bilateral trade flows at the exporting country x importing country x HS6 product x year level.

[11] Commuting data ignores both commuting within the city (which is rare in the Paris case), as well as house movers in the same city. On the contrary, commuting data at the city level maybe overstate the change in commuting for workers who changed residence and of city and are located far from the centroid of the city.

[12] Inside the EU, firms are required to report their shipments by product and destination country only if their annual export value exceeds the threshold of 150,000 Euros. For exports outside the EU, all flows are recorded unless their value is smaller than 1,000 Euros or one ton. Yet, these thresholds eliminate a very small share of the total exports (see Berman, Berthou, and Héricourt (2015)) .

Final samples and descriptive statistics.

Our data allow us to estimate the impact of wage variations on the worker-level home-workplace distance change in both the short- and the long run. To do so, we use two distinct samples. First, the short-run estimates are derived from the full sample of exporting firms' employees. This sample contains a bit less than 2 billion observations covering around 30000 firms. Table 1-A displays the descriptive statistics of this full sample.

Second, to focus on the long-run elasticity, we use another sample which is composed of exporters' employees that, indeed, have moved out of their houses. From the full sample, we can identify workers who effectively changed their residence cities. We consider these workers as "movers" and these constitute our reduced sample. As these observations exclude moving frictions, we consider that they allow us to estimate the long-run elasticity (as in Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016)). We thus drop workers for which change in commute time is zero as they remained in the same residence city. Table 1 – B shows the descriptive statistics of that reduced sample. There are no large differences in the composition of these two samples.

The external validity of our result could be threatened, as we focus on exporting firms' employment only. Yet, this sample selection is likely unrelated to commuting behavior. In particular, exporting firms' workers may have particular characteristics in terms of skills and productivity that are different from those of domestic-only firms' workers. In many contexts, choosing to "isolate" exporting firms' workers leads to many methodological issues both theoretically and empirically. However, to the best of our knowledge, there is no reason or existing study that documents that the individual location decision depends on the type of firm the worker is employed in. The crucial parameter here is rather income[13]. On top of that, if there were any bias across workers in this dimension, our exercise includes individual worker fixed effects, reducing this concern.

4. Empirical strategy

4.1 General Relationship

We identify the magnitude of the commuting distance elasticity to wage by estimating the following general equation:

$$ln\, Dist_{ift} = \alpha_1 ln\, W_{ift} + \alpha_2 C_{ft} + \alpha_3 C_{it} + FE + \varepsilon^{OLS}_{ift}$$

where $Dist_{ift}$ is the commuting distance of any worker i in firm f in year t. The main dependent variable is the log yearly wage of that worker, W_{ift}. Parameter α_1 captures the elasticity of commuting distance to wage and is of prime interest to us. We include a set of firm-level controls,

[13] Indeed, a worker employed by an exporting firm would have the same preferences regarding housing consumption, commuting and amenities as a worker employed by a local firm, as far as they have the same wages.

C_{ft} (log Assets, Log Apparent Labor Productivity, defined as total value added per worker and log employment of the plant), and a set of worker-level controls, C_{it} (i.e. individual observable worker characteristics: skilled/unskilled dummy, occupation group, age, gender).

We include a set of fixed effects, $\boldsymbol{FE}$, to account for unobserved heterogeneity across observations. We crucially include a worker fixed effect (i) to account for unobserved heterogeneity across workers and (ii) to properly identify our effect. When a worker FE is included, α captures the within-worker effect of earnings on commuting distance. In other words, when we include this FE, we ensure that we estimate a within-worker effect instead of any across-workers effect. The identifying variation is for any worker across years, thus absorbing potential variations across workers. Among others, worker self-selection into specific firms or differences in employment structures across firms are all excluded from the analysis. As a counterpart, identification only arises from wage (and commuting distance) variation across years. Since we follow workers only for a two-year period, our estimates capture *immediate* changes in location. We also include a year fixed effect, to account for unobserved heterogeneity across years (absorbing economy-wide shocks), and in some specifications, an occupation-specific fixed effect. Finally, ε_{ift}^{OLS} is a random error term capturing all omitted factors, which we allow to be heteroskedastic and correlated across HS2 sectors: in practice, we report the standard errors clustered at the HS2 sector-year level, because the trade shocks mainly occur at this aggregate level.

Our setup yields a difference-in-difference analysis. Indeed, the empirical model includes a worker-specific (that subsumes the employing firm) and a year fixed effects. The estimates inform on the impact of changes in wages (at the worker-year level) on changes in the commuting distance (also at the worker-year level), controlling for worker and year unobserved shocks. The estimates are thus obtained by comparing variations in the commuting distance (over time) across workers, which are differently affected by shocks.

We estimate this equation on both the full and the reduced samples to identify the short-run and long-run elasticities, respectively. Intuitively, we always regress the commuting distance on earnings, conditional on worker fixed effects. This worker fixed effect implies we regress the change in the commuting distance on the changes in earnings, where changes can be zero. The short-run elasticity is estimated using the full sample, whereas the long-run model excludes workers for which the change in commuting distance is zero (because they did not change residence city).

4.2 Endogeneity concerns

Independently on the sample, we could estimate this equation with an OLS estimator but there are many reasons to believe that these estimates are biased. First, firms may compensate workers for their commuting costs. Even though wage data does not include monetary compensation for the commuting costs, we cannot exclude that wages include compensation as a wage package bargained when hired.[14] Second, workers may simultaneously decide where to live and where to work[15], leading to simul Third, other determinants of the locational decision are absent from our exercise. Among others, we have no information about the marital status of the worker nor on the family composition. We also can't control whether the worker owns her home or if she rents it. This raises some bias related to omitted variables. The worker FE partially controls for this issue, to the extent these variables don't change over time, within worker.

We thus have to deal with endogeneity issues due to simultaneity and omitted variables. Unobserved variables may affect both wages and commute, causing spurious correlations between these variables. The OLS coefficient could thus be a biased estimator of the true parameter. To overcome endogeneity and to allow for a causal interpretation of the coefficients, we use an instrumental variable (IV) strategy in a two-stage least square (2SLS) estimation. Our IV strategy has to identify sources of variations in wages that are exogenous and uncorrelated to the firm and to workers, thus excluding the simultaneous variation in wages and distance. Identification of this relationship requires instruments that (i) are related to wages and not to location or commuting decisions (except through wages), and (ii) that are orthogonal to firm-specific decisions. We advocate that *shifters* of exports at the firm level satisfy these two conditions.

4.3 Instrumental Variable and 2SLS estimation

On the one hand, recent evidence supports that firm-level exports affect wages at the worker level. Amiti and Davis (2012) show that a decline in output tariffs raises the wages of workers at firms that export. In the French case, the positive impact of trade on wages is supported by recent findings (see Carluccio, Fougère, and Gautier (2015)). In this respect, exports are well connected to wages. It is also hard to argue that exports are related to commuting patterns at the individual level. With this in mind, we will thus use a *shifter* of firm-level exports as an instrument that, by construction, is orthogonal to firm-specific supply choices.

We measure trade shocks using a world import demand (*WID* hereafter) addressed to the products that are sold by the firm, so as to capture exogenous changes in trade conditions. Our baseline trade shock measure is constructed using information about the foreign demand addressed to the firm using product and destination information. Specifically, we compute the sum of world imports in the products-destinations in year t (using the BACI dataset and excluding France as an exporter) weighted by the share of each product-destination in the firm total exports in that year t (using the firm-level exports data). Weights are computed using the yearly share of the product-destination in the firms' total exports. A product is defined at the 6-digit (HS6) level. More precisely, we define

[14] Note that in France, wages are rarely bargained, or at a low scale, and it is all the less likely during that period of time, with high unemployment in France and in the Ile-de-France.
[15] This possibility is only partly ruled out in our sample when excluding job movers.

$$WID_{ft} = \sum_{js} \omega_{fjst} \times M_{jst}$$

where ω_{fjst} is the share of each product *s* and destination *j* in firm *f* exports in year *t*, M_{jst} is the total value of imports for product *s* and destination *j* in year *t*, excluding France as potential exporter. By excluding French exports to this destination, we exclude sources of variations that originate in France and may be correlated with changes in the firm.

We settle on extensive literature using trade shocks as shifters in empirical exercises. This measure is related to a shift-share instrument following the seminal contribution of Bartik (1991). The use of this type of variable, combining a firm-level exposure (the "share") to a set of exogenous shocks (the "shifts") has been recently increasing and applied in many contexts.[16] This type of measure has in particular been widely used on French firm-level trade data (i.e. on the same data as we use here) to isolate trade shifters, which are independent of firms' and workers' choices (Mayer, Melitz, and Ottaviano (2014); Hummels et al. (2014); Berman, Berthou, and Héricourt (2015); Aghion et al. (2018)). Interestingly, Aghion et al. (2018) show that trade shifters are uncorrelated to many firm-level outcomes.

We include this instrument -- WID_{ft} -- explicitly in the estimation. The following equations assess the effect of exogenous changes in wages (occurring through variations in the instruments) on commuting distance, controlling for firm characteristics:

$$\ln W_{ift} = \beta_1 \ln WID_{ft} + \beta_2 C_{ft} + \beta_3 C_{it} + FE + \varepsilon_{ift}^{first-stage}$$

$$\ln Dist_{ift} = \alpha_1 ln \hat{W}_{ift} + \alpha_2 C_{ft} + \alpha_3 C_{it} + FE + \varepsilon_{ift}^{second-stage}$$

where $ln \hat{W}_{ift}$ is the predicted value of the log wage from the first-stage equation. This strategy thus uses wage variations, coming from trade shifts, on commuting patterns. We estimate these equations using two-stage least-squares, and standard errors are still clustered at the HS2 – year level.

[16] Among others, Autor et al. (2013) use a shift-share IV to investigate the consequences of the "China shock" on labor markets. Peri, Shih, and Sparber (2015) focus on the impact of foreign workers on native's wages while Imbert et al. (2018) use a shift-share IV strategy to estimate the causal effect of rural-urban migration on production in China. Bombardini and Li (2020) study the impact of trade expansion on local pollution and health outcomes in China too. Note also that the inference properties of this technique have been recently extensively discussed in Borusyak and Hull (2020); Borusyak, Hull, and Jaravel (2018); Adao, Kolesár, and Morales (2019); Goldsmith-Pinkham, Sorkin, and Swift (2020).

5. Results

5.1 Baseline results in the short-run

IV Results.

Our preferred estimation is obtained using an IV-2SLS estimation in which we instrument the (log) yearly wage by the (log) exogenous trade shifter. In the first stage, we estimate the impact of trade shocks on individual wages, controlling for firm and workers' characteristics. In the second stage, we use the predicted value of wage from the first stage and estimate its impact on commuting distance. Since world import demand is uncorrelated to the firm and to workers, the estimates we obtain in the second stage are the causal impact of changes in wages on distance. In other words, we measure how much commuting distance is affected by variations in wages coming from exogenous trade shifters.

Table 2 presents the IV-2SLS results from the two stages. Each column corresponds to a different specification. The upper panel of Table 2 presents the first-stage results, and the bottom panel presents the second-stage, main results. Columns (1) only include worker and year FE. Column (2) adds firm-year controls (log Assets, log Apparent Labor Productivity and log Employment) while column (3) includes worker-level controls (such as age, gender, skilled position). In column (4), we include both sets of controls and in column (5) we add a set of 2-digit occupation fixed effects. Specifications (4) and (5) are our preferred estimations.

The top panel reports the first-stage results. Across specifications, coefficients all stand around 0.015 and precision is quite high. Since the specification is log-log, we estimate that a 10% increase in world import demand is associated with a 0.15% increase in wage, ceteris paribus. This result is robust across specifications and in line with existing evidence regarding the pro-wage effects of increased export opportunities. We always estimate that changes in world demand are positively and significantly associated with changes in wages.

The bottom panel reports the second-step estimation results and we only report the coefficients on the log wage, which is of prime interest to us. We always estimate a positive and significant coefficient to wage on commuting distance. Our preferred specification is in column (4): in this demanding estimation, results are consistent with a positive impact of variation in wages on commuting distance for a given worker across years. Our estimates mainly lie around 0.2. We estimate that a 10% increase in wage is associated with at least a 1.8% increase in commuting distance on average.

Quantitatively, our estimates are pretty small but can be explained by three reasons. First, we estimate contemporaneous changes in distance associated with changes in wage. For sure, there are many barriers leading to a slow and sticky change in commuting patterns, inducing a downward bias in the coefficient. Second, our data does not cover the total income of households, which has been documented to be the main determinant of residence change in couples and families. By approximating the household income with individual wages, we identify the consequences of the increase in wages for *at least* one person in the household. Third, we only focus on the intensive margin of change. The true estimate of the distance elasticity should in a sense include job stayers (our sample) and job movers (which are not covered in the raw data). Our estimates are consequently likely to be negatively biased, as these job movers may exhibit large changes in wages and thus in distance.

OLS results.

For comparison, Table 3 displays the OLS estimation results. All specifications display a positive and significant coefficient, in line with the IV estimates above. Yet, the magnitudes of the coefficient are different from the IV-2SLS results. Contrasting with the IV results, OLS coefficients are hardly economically meaningful. Before controlling for the endogeneity issues between wages and distance, we would estimate a close to null relationship between wages and commuting distance. Armed with our IV and with OLS results, we infer that (i) endogeneity concerns are real and harmful to the identification of the elasticity and (ii) that the relationship is likely to be from wages to distance.

5.2 Baseline results in the long-run

Table 4 presents the long-run IV results. Table 4 replicates Table 2 using the "movers" sample. Doing so allows us to identify the long-run estimates, as we abstract from moving frictions by focusing on households that effectively changed residence. First-stage results are very close, in coefficient magnitude, to the results on the full sample. This is reassuring as this first-stage estimation (using shifts in wages coming from trade activities) is not supposed to be different across households or across samples. Second-stage results are however different as we estimate larger point estimates, compared to Table 3. In the long run, the commuting distance elasticity to wages is larger (around twice as large) than in the short run[17]. The benchmark estimate is around 0.3,

[17] Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016) also compare estimates on two samples, either including residential movers or not. They find that the short-run estimate is "about three times smaller" than the medium-/long-run estimate in their exercise. See their Table 2, column 1.

suggesting that the same 10% increase in wages increases commuting distance by 3% in the long run and around 1.8% in the short run (column 4, Tables 2 and 4). Compared to the whole sample, estimates in the long run are less precisely estimated: standard errors are about 3 times larger in that sample, which is driven by the decrease in sample size.

Table 5 shows the OLS results in the long run. All specifications display a positive and significant coefficient, in line with the IV estimates above. Yet, once again, comparing the OLS and IV confirms the importance of accounting for endogeneity concerns.

5.3 Effects on commuting patterns

We replicate the same type of estimation but focus on alternative dependent variables, all related to the location decisions of workers. In particular, since our strategy relies on a diff-in-diff estimation, our data allows us to study the location choice of the workers. By including information about *city* characteristics in our dataset, we can interpret the choice of workers in terms of economic characteristics of new home cities compared to former home cities.

Table 6 provides the second-stage estimation of the following equation:

$$CityCharact_{ifct} = \alpha_1 ln\, \hat{W}_{ift} + \alpha_2 C_{ft} + \alpha_3 C_{it} + FE + \varepsilon_{ift}$$

where $CityCharact_{ifct}$ represents city c characteristics -- such as distance from the historical center, employment, density, revenues... -- in which worker i employed in firm f lives. By comparing city characteristics across years for a given worker, the coefficient α_1 captures whether a change in wage correlates with differences in home city characteristics. Panel A provides the short-run estimates whereas Panel B displays the long-run results.

We consider many city characteristics in our exercise. Column (1) focuses on the impact of changes in wages on the distance to the center of Paris. We compute this distance as the distance with respect to the 75001 postal code, which corresponds to the first *arrondissement* of the city of Paris, which is a proxy of the "historical" center of Paris. We thus obtain a set of unilateral distances at the city level, and the coefficient is identified from the difference across years in the distance between central Paris and the home city for any worker. Previous results supported that exogenous changes in wages led workers to move away from their firms. The negative and significant coefficient in column (1) suggests that, on top of this, workers moved closer to the center of Paris, i.e. towards the historic center of Paris. This result is important as we obtain a similar qualitative result with respect to Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016).[18] Workers that move tend to go

[18] This is the same mechanism as the one suggested in Gutiérrez-i-Puigarnau, Mulalic, and van Ommeren (2016), yet in their cases, as jobs are also located in the historical center, commuting distance decreases with wage increases. In the Paris case,

towards the historic center of Paris. One plausible explanation is that in the Paris Region, the center has a strong amenity advantage over the suburbs. Indeed, the "exogenous" natural and historical amenities exert a power of attraction on the households that can afford housing in their surroundings (see recent evidence in Gaigné et al. (2022)).

Column (2) estimates the correlation between changes in wages and the average revenue of cities in 2002, which is out of our sample. We use the full DADS dataset (see the Data section) and compute the average yearly labor revenue for each (home) city. We estimate that on average, workers that experienced increases in wages tend to move to richer cities, compared to their former home city. This is consistent in a sense with Hedman et al. (2011), Bailey and Livingston (2007), Quillian (2003) and Boterman (2012). Indeed, these papers exhibit the fact that when moving, households tend to choose to settle in areas where the demographic and socio-economic characteristics of the inhabitants are close to theirs. Thus, if a rise in income leads to a move, the new neighborhood is very likely to be "richer" than the old one.

Does density matter for the location choice? We measure worker, firm, and exporting density as the total number of workers, firms (both from the DADS dataset), and exporting firms (from the Customs dataset) per square kilometer in each city in 2002. City size data comes from the CORINE Land Cover (CLC) dataset. Columns (3) to (5) state that workers that experienced increases in wages tend to move to cities with higher workers' density (column (3)), but lower firms' density (column (4)) as well as lower exporting firms' density (column (5)), compared to their former home city.

Moving towards rich cities, where house prices are higher, where jobs are relatively scarce but urban amenities abundant when income rises rules-out findings stating that rich households would naturally pay for shorter commutes or for larger houses. In this, our results inform on the crucial role of amenities. Richer workers increase their commuting but *not* for lower housing prices, neither for shorter home-workplace trips. The main motivation here seems to be the amenities. This echoes recent findings emphasizing the role of urban amenities (see section 2.1).

Concerning commuting modes[19], by comparing the potential travel times by public transport and by private car (see Data section), our model is able to provide some insights regarding the evolution (for a given worker) of the relative accessibility of workplace to home by public transport. At the worker level, given home and work cities, we can measure the relative time cost of using cars compared to public transport and its evolution (given residency changes). Table 6, column (6)

---

jobs are not located in the historical center: when workers move towards the historical center to live, they get further away from their jobs and closer to urban amenities.

[19] The limit of the exercise is that we do not know the commuting mode of workers.

shows that increases in wages are associated with a higher relative time cost of using cars (potentially increasing public transport use). Richer people are moving in space and the new home-work combinations are more consistent with the use of public transportation (than the previous home-work combinations). This result should however be looked at with caution. Specifically, it does not state (here in the case of Paris Region) that the public transportation system performance in the new home city determines the location choice of the households. It only states that the comparative advantage of the car compared to public transport for commuting is lower in the new city of residence. However, the new cities of residence tend to be more central than the old ones, and the center of Paris has historically been better equipped with public transport than the suburbs. That being said, residing in areas where the comparative advantage of the car is lower could encourage some workers to trade their cars for public transport.

Long-run results confirm the sign and significance of the results on the whole sample. We obtain larger point estimates, in line with previous outcomes, and with intuition as these moving workers are identifying the coefficients.

6. Robustness checks

6.1 Omitted Variables

Worker-level Wages and Household Earnings.

Numerous works in urban economics show that it is the overall household income that affects residential location, and not only labor income. Our data however does not allow us to know the household composition and how total household earnings evolve over time. To overcome this bias, we introduce a set of additional fixed effects in the analysis to (implicitly) capture the dynamics of wage in the neighboring of the worker.

Table 7 provides the results with more demanding fixed effects. We introduce (i) a home-city fixed effect (N=1296 i.e. for each of the 1296 "cities" in our setup), (ii) a home-city time-varying fixed effect (also capturing time-varying amenities for instance), and (iii) a home-city time-varying sector-specific fixed effects, that all capture the heterogeneous dynamics of wages in various dimensions, and that could account for the average wage dynamics of the household. Results show that accounting for these fixed effects confirms the main result and tends to increase the value of the point estimate, compared to the benchmark.

Agglomeration and co-agglomeration.

On average, firms in the same or close industries will tend to concentrate in the same local labor markets. As a result, some local labor markets will systematically experience growth from a set of related firm shocks, and these same local labor markets will experience increases in housing prices and wages. We account for this omitted variable bias by including additional fixed effects, grouping work cities into local labor markets (N=26) (then denoted EZ for "employment zones", "*zones d'emploi*" in French) using nationwide classification. Columns (4) and (5) of Table 7 include EZ fixed effects and EZ-year fixed effects, respectively. Point estimates are smaller in magnitude compared to the baseline effects, but significance remains. We interpret this as a confirmation that broad economic shocks and local market effects could drive part of the estimated effect.

Industry trends.

An implicit assumption in the analysis is that workers' sorting into industries is not related to industry specificities (that could be correlated with trade shocks or trade growth at the industry level). Whereas we cannot rule this assumption out (as we do not observe workers' sorting), we introduce some controls for trends in the analysis. For each firm-year, we identify the specific 4-digit industry of the main exported product (N=1254) and include a set of industry x year linear trends in the general specification. Results are presented in column (6) of Table 7. When accounting for disaggregated industry-specific linear trends, results are not quantitatively affected.

Domestic Sales.

We also checked that our world import demand does not capture business cycles conditions that may also affect domestic demand. World business cycles conditions may affect both foreign sales and domestic sales. We thus insulate our results from domestic sales. We include domestic sales, from the balance-sheet data, in both stages. The remaining effect thus identifies the change in distance controlling for the simultaneous change in domestic sales. Note that we include this variable in both stages. In the first stage, we recover the prediction of wages, now controlling for domestic sales.

Tables A1 and A2 in the appendix provide the results of the IV regressions controlling for domestic demand. Our baseline results are not affected by this additional control variable. The magnitudes of the coefficients remain unchanged compared to the baseline coefficient.

Previous Export Sales.

The difference between OLS and IV results could be driven by the inclusion of export-related variables in the exercise, through the instrument. Exports sales are for instance absent from the baseline estimation, but, de facto, included in the instrument and thus in the IV results. We checked that this omission does not affect the results. Tables A3 and A4 in the appendix show the results when controlling for lagged exports outcomes, on top of other firm-year variables. We obtain close results in magnitude and in precision. We are thus confident that the IV strategy does not alter the results through the inclusion of export-related variables.

6.2 Sensitivity checks

We checked the sensitivity of the main parameter in our IV estimations when using alternative measures of the main variables. Instead of providing an extensive set of tables, we follow Simonsohn, Simmons, and Nelson (2020) and provide a "specification curve" regarding the coefficient of interest (α_1). We replicate variants of the IV specification, with the complete set of fixed effects, and alternatively replace the dependent variable (the measure of distance), the measure of wage, and the weighting scheme in the world import demand as the instrument. In formal terms, we estimate variants of the following system of equations:

$$\ln Wage_{ift} = \beta_1 \ln Instrument_{ft} + \beta_2 C_{ft} + \beta_3 C_{it} + FE + \varepsilon_{ift}^{first-stage}$$

$$\ln Commuting_{ift} = \alpha_1 ln\, \widehat{Wage}_{ift} + \alpha_2 C_{ft} + \alpha_3 C_{it} + FE + \varepsilon_{ift}^{second-stage}$$

From these estimations, we obtain a set of estimates of α_1 (N=240) that we then plot in Figure 1 and Figure 2[20]. Figure 1 shows the results for the full sample, i.e. the short-run estimates, whereas Figure 2 displays the long-run estimates. In each Figure, the top graph shows all estimated wage coefficients, with the associated confidence intervals (using clustered standard errors, at the sector-year level, as in the benchmark estimations) and the bottom shows the characteristics of the estimation from which the coefficient is obtained.

Alternative commuting measures.

Our baseline measure of commuting was the commuting distance using public transport. As alternatives, we consider (i) commuting distance using a private car, as well as (ii) commuting time using public transportation and (iii) private car. We do this to check that results hold independently of the measure, all the more that we do not know the transportation mode that is used at the individual level. Estimates appear independent to the measure we use and thus hold whatever the transportation mode. If any effect occurs, point estimates suggest that elasticity of travel time by public transport seems to be the least elastic measure to changes in wages, while the elasticity of travel distance in private cars is estimated to be larger than the other modes.

Alternative wages measures.

Second, Figures 1 and 2 show the sensitivity of the results to the use of alternative wage measures. While our baseline results were based on yearly net wage (i.e. the yearly disposable individual labor income, net of social contributions but not of income tax), results hold when considering (i) yearly

[20] In each Figure, we combine each of the 4 alternative measures of commuting, with the 4 measures of wages, and 5 weighting schemes for the instrument and 3 measures of trade shocks, leading us to 240 combinations. More details in the next paragraphs.

gross wage (including social contributions) as well as (ii) hourly net and (iii) hourly gross wage. Elasticities of commuting to hourly wage appear larger (point estimates) but are less precisely estimated. Overall, we do not estimate meaningful differences across wage measures.

Alternative weights in world import demands.

Third, we check the sensitivity of our results when using an alternative weighting scheme in our measure of world import demand. Our baseline measure used the effective share of each product-destination in the firm's total exports. We first consider the sum of world imports in the products-destinations served by the firm in year t weighted by the *average* share of each product-destination in the firm's total exports:

$$WID_{ft}^{average} = \sum_{js} \omega_{fjs} \times M_{jst}$$

where ω_{fjs} is the average share of each product s and destination j in firm f exports over the total period. Second, we also computed a world import demand using product-destination *lagged* shares in exports as a weighting scheme:

$$WID_{ft}^{1yr-lag} = \sum_{js} \omega_{fjs,t-1} \times M_{jst}$$

where $\omega_{fjs,t-1}$ is the 1-year lagged share of the product-destination in total firms exports. Both measures check that time-variation of the foreign demand measure comes from the country-level imports by product, not from the firm-level weights. We also consider further lags in time to reduce concerns, using 2-year lags ($WID_{ft}^{2yr-lag} = \sum_{js} \omega_{fjs,t-2} \times M_{jst}$) and 5-year lags ($WID_{ft}^{5yr-lag} = \sum_{js} \omega_{fjs,t-5} \times M_{jst}$).

Figure 1 shows that the positive elasticity of commuting to wages holds independently of the weighting scheme of foreign shocks and of the firm exposure to export demand shocks. All estimates are positive and significant. Second-stage point estimates are very close to our baseline results using both average and lagged weights while firm exposure to export shocks does not affect our results.

Alternative foreign shocks measures.

Fourth, we check the sensitivity of our results when considering growth in foreign product-level demand shocks instead of their levels. Many works in the literature (following Autor, Dorn, and Hanson (2013) for instance) use 5-year or 10-year changes in shocks as the identifying variation, instead of the contemporaneous levels. We follow this intuition and compute the alternative measures of demand shocks, focusing on their 5-year growth:

$$WID_{ft}^{5y-growth} = \sum_{js} \omega_{fjst} \times \Delta_5 M_{jst}$$

in which $\Delta_5 M_{jst} = M_{jst} - M_{js,t-5}$ denotes the 5-year lag in country-product specific shifts in imports. We also consider for additional robustness a 10-year change in demand shocks:

$$WID_{ft}^{10y-growth} = \sum_{js} \omega_{fjst} \times \Delta_{10} M_{jst}.$$

Combining the weighting schemes (the above equations use the baseline weighting scheme for illustration) with each of the three foreign shocks measures (contemporaneous, 5-year and 10-year change) provides 15 alternative instruments. Figure 1 shows that our estimates hold when we use changes in shocks instead of the shocks in levels in the IV estimation. We obtain very close estimates independently of the measure we use (in levels or in change). We however obtain some non-significant estimates only when using changes in demand shocks with lagged shares as a weighting scheme.

Overall, Figures 1 and 2 display a robust pattern regarding the main estimates. Independently of the measures used in the estimation, we always estimate a positive coefficient, which is in the vast majority significant. Short-run estimates lie in the 0.1 to 0.3 range (which is consistent with the baseline estimated elasticity of 0.15) whereas we estimate a larger long-run elasticity, around 0.37 on average across specifications.

6.3 Pre-trends and heterogeneous effects over time concerns

Our analysis could be biased by standard problems in difference-in-difference analyses, including pre-trends and heterogeneous effects over time. Indeed, these issues are likely to arise in standard diff-in-diff analyses and in particular in those using shift-share instruments (Goldsmith-Pinkham, Sorkin, and Swift (2020)). We follow the tests proposed in Goldsmith-Pinkham, Sorkin, and Swift (2020) and test the assumption of parallel trends. In our context, firms (establishments) are exposed to foreign shocks in each period. We check whether the commuting distance of workers is affected by the shocks in *t* and we allow for the lagged and forward impact of the shocks from *t-4* to *t+4*. We estimate :

$$ln\, Dist_{ift} = \sum_{\tau=-4}^{4} \zeta_\tau\, ln\, WID_{f,t+\tau} + \lambda_i + \lambda_t + \varepsilon_{ift}$$

where the ζ_τ are the lagged and forward effects of import demands on commuting distance. Figure 3 plots the estimated ζ_τ, using the benchmark WID measure.

Results confirm the parallel trend assumption. We estimate that “treated” and “untreated” workers in *t* do not have any differential commuting distance before the treatment. The effect of import demand shocks are concentrated around the year of treatment and in *t+1*. We estimate only marginal effects of forward import demand shocks on commuting distance. Overall, we do not estimate strong heterogeneous effects across time, except in *t+1*, confirming the absence of anticipation of workers. Appendix C displays further results regarding (i) an augmented specification of the above equation including interactions between initial controls and year fixed effects and (ii) the other measures of import demand (see the discussion about weights and shocks in subsection 6.2).

6.4 Plausibly exogenous IV

The validity of our IV results relies on the validity of the excluded instrument we used. Appendix B checks the validity of the world import demand as an instrument and assesses the sensitivity of the second-stage results. Briefly, (i) there is no correlation between our excluded instrument and firms' and workers' observables, (ii) controlling for the average (and potentially non-random) exposure of firms and workers to random foreign shocks – following Borusyak and Hull (2020)—, we obtain close estimates to our baseline results, (iii) IV results appear stable despite potential minor violations of the exclusion restriction (Conley, Hansen, and Rossi (2012)) and (iv) using a second excluded instrument, our results satisfy standard tests for overidentifying restrictions.

7. Conclusions

We estimate a causal positive worker-level commuting distance elasticity to wage, using datasets covering the universe of Ile-de-France (surrounding Paris). To estimate this elasticity, we use export shocks at the firm level as a wage shifter at the individual level. We identified trade shocks that are uncorrelated to firms' and workers' decisions by computing, using disaggregated customs data, the world import demand that is addressed to the firm. On average, workers increase their commuting distance when wages exogenously increase and tend to get closer to the historic city center. Overall, we estimate a 0.15 elasticity in the short run and a 0.35 elasticity in the long run.

This paper provides clear identification of the wage-commuting relationship. Our result also supports that the impact of wage variations on city structures depends both upon urban amenities and on jobs locations. In our particular geographical framework, jobs are mainly located outside the city's historical center—whereas amenities are in the city—, while it may not be the case in other cities. Understanding how the geography of jobs shapes the city structure calls for future research. Our paper also calls for future research on the consequences of trade activities, beyond their effect on wages, on workers' location within and around cities. Results in this area are all the more necessary that expanding trade is likely to generate many unexplored land use and urban planning issues.

References

ADAO, R., M. KOLESÁR, AND E. MORALES. (2019): "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 134, 1949–2010.

AGHION, P., A. BERGEAUD, M. LEQUIEN, AND M. J. MELITZ. (2018): The Heterogeneous Impact of Market Size on Innovation: Evidence from French Firm-Level Exports,. Working Paper Series.

ALONSO, W. (1964): *Location and Land Use*, Cambridge: Harvard University Press.

AMITI, M., AND D. R. DAVIS. (2012): "Trade, Firms, and Wages: Theory and Evidence," *Review of Economic Studies*, 79, 1–36.

ARNOTT, R. (1998): "Economic theory and the spatial mismatch hypothesis," *Urban Studies*, 35, 1171–85.

AUTOR, D., D. DORN, AND G. HANSON. (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–68.

AUTOR, D., D. DORN, G. H. HANSON, G. PISANO, AND P. SHU. (2020): "Foreign Competition and Domestic Innovation: Evidence from US Patents," *American Economic Review: Insights*, 2, 357–74.

Bailey, N., and Livingston, M. (2007) : "Population turnover and area deprivation", Policy Press.

BARTIK, T. J. (1991): "Who benefits from state and local economic development policies?,"

BARTOLOME, C. DE, AND S. ROSS. (2003): "Equilibria with local governments and commuting: income sorting vs income mixing," *Journal of Urban Economics*, 54, 1–20.

BERMAN, N., A. BERTHOU, AND J. HÉRICOURT. (2015): "Export dynamics and sales at home," *Journal of International Economics*, 96, 298–310.

BOMBARDINI, M., AND B. LI. (2020): "Trade, pollution and mortality in china," *Journal of International Economics*, , 103321.

BORUSYAK, K., AND P. HULL. (2020): Non-Random Exposure to Exogenous Shocks: Theory and Applications,. Working Paper Series.

BORUSYAK, K., P. HULL, AND X. JARAVEL. (2018): Quasi-Experimental Shift-Share Research Designs,.

Boterman, W. R. (2012): "Residential mobility of urban middle classes in the field of parenthood", Environment and Planning A: Economy and Space, 44(10), 2397–2412.

BRUECKNER, J. K., AND S. S. ROSENTHAL. (2009): "Gentrification and neighborhood housing cycles: Will America's future downtowns be rich?," *The Review of Economics and Statistics*, 91, 725–43.

BRUECKNER, J. K., J.-F. THISSE, AND Y. ZENOU. (1999): "Why is central Paris rich and downtown Detroit poor?: An amenity-based theory," *European economic review*, 43, 91–107.

CARLUCCIO, J., D. FOUGÈRE, AND E. GAUTIER. (2015): "Trade, Wages and Collective Bargaining: Evidence from France," *The Economic Journal*, 125, 803–37.

CONLEY, T. G., C. B. HANSEN, AND P. E. ROSSI. (2012): "Plausibly exogenous," *Review of Economics and Statistics*, 94, 260–72.

COUTURE, V., C. GAUBERT, J. HANDBURY, AND E. HURST. (2019): "Income Growth and the Distributional

Effects of Urban Spatial Sorting," *National Bureau of Economic Research Working Paper Series*, No. 26142.

Couture, V., and J. Handbury. (2017): "Urban Revival in America, 2000 to 2010," *National Bureau of Economic Research Working Paper Series*, No. 24084.

Cuberes, D., J. Roberts, and C. Sechel. (2019): "Household location in English cities," *Regional Science and Urban Economics*, 75, 120–35.

Duranton, G., and D. Puga. (2014): "The Growth of Cities," in *Handbook of Economic Growth*, Elsevier, 781–853.

Fu, S., and S. L. Ross. (2013): "Wage Premia in Employment Clusters: How Important Is Worker Heterogeneity?," *Journal of Labor Economics*, 31. Journal of Labor Economics, 271–304.

Fujita, M. (1989): *Urban Economic Theory*, . Cambridge BooksCambridge University Press.

Fujita, M., J.-F. Thisse, and Y. Zenou. (1997): "On the Endogeneous Formation of Secondary Employment Centers in a City," *Journal of Urban Economics*, 41, 337–57.

Gaigné, C., H. R. A. Koster, F. Moizeau, and J.-F. Thisse. (2022): "Who lives where in the city? Amenities, commuting and income sorting," *Journal of Urban Economics*, 128, 103394.

Gaulier, G., and S. Zignago. (2010): "BACI: International Trade Database at the Product-Level. The 1994-2007 Version," *CEPII Working Paper*, .

Giménez-Nadal, J. I., J. A. Molina, and J. Velilla. (2018): "Spatial distribution of US employment in an urban efficiency wage setting," *Journal of Regional Science*, 58, 141–58.

Glaeser, E. L., M. E. Kahn, and J. Rappaport. (2008): "Why do the poor live in cities? The role of public transportation," *Journal of Urban Economics*, 63, 1–24.

Glaeser, E. L., J. Kolko, and A. Saiz. (2001): "Consumer city," *Journal of Economic Geography*, 1, 27–50.

Gobillon, L., H. Selod, and Y. Zenou. (2007): "The mechanisms of spatial mismatch," *Urban Studies*, 44, 2401–28.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift. (2020): "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 110, 2586–2624.

Guerrieri, V., D. Hartley, and E. Hurst. (2013): "Endogenous gentrification and housing price dynamics," *Journal of Public Economics*, 100, 45–60.

Gutiérrez-i-Puigarnau, E., I. Mulalic, and J. van Ommeren. (2016): "Do rich households live farther away from their workplaces?," *Journal of Economic Geography*, 16, 177–201.

Gyourko, J., C. Mayer, and T. Sinai. (2013): "Superstar Cities," *American Economic Journal: Economic Policy*, 5, 167–99.

Hedman, L., M. van Ham, and D. Manley. (2011): "Neighbourhood Choice and Neighbourhood Reproduction," *Environment and Planning A: Economy and Space*, 43, 1381–99.

Hoelzlein, M. (2019): "Two-Sided Sorting and Spatial Inequality in Cities," *Mimeo*, .

Hummels, D., R. Jørgensen, J. Munch, and C. Xiang. (2014): "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review*, 104, 1597–1629.

Imbert, C., M. Seror, Y. Zhang, and Y. Zylberberg. (2018): "Migrants and firms: Evidence from china," *Mimeo*, .

JIN, J., AND K. PAULSEN. (2018): "Does accessibility matter? Understanding the effect of job accessibility on labour market outcomes," *Urban Studies*, 55, 91–115.

KAIN, J. F. (1968): "Housing Segregation, Negro Employment, and Metropolitan Decentralization," *The Quarterly Journal of Economics*, 82, 175–97.

KOSTER, H. R. A., J. N. VAN OMMEREN, AND P. RIETVELD. (2014): "Historic amenities, income and sorting of households," *Journal of Economic Geography*, 16, 203–236–203–236.

MANNING, A. (2003): "The real thin theory: monopsony in modern labour markets," *Labour Economics*, 10, 105–31.

MAYER, T., M. J. MELITZ, AND G. I. P. OTTAVIANO. (2014): "Market size, competition, and the product mix of exporters," *American Economic Review*, 104, 495–536.

---. (2021): "Product Mix and Firm Productivity Responses to Trade Competition," *The Review of Economics and Statistics*, 103, 874–91.

MCKINNISH, T., R. WALSH, AND T. KIRK WHITE. (2010): "Who gentrifies low-income neighborhoods?," *Journal of Urban Economics*, 67, 180–93.

MILLS, E. S. (1972): *Studies in the Structure of the Urban Economy*, Baltimore: The Johns Hopkins Press.

MONTE, F., S. J. REDDING, AND E. ROSSI-HANSBERG. (2018): "Commuting, Migration, and Local Employment Elasticities," *American Economic Review*, 108, 3855–90.

MORETTI, E. (2011): "Local Labor Markets," in *Handbook of Labor Economics*, ed. by Ashenfelter, O. and D. Card. Elsevier, 1237–1313.

MULALIC, I., J. N. VAN OMMEREN, AND N. PILEGAARD. (2014): "Wages and Commuting: Quasi-natural Experiments' Evidence from Firms that Relocate," *Economic Journal*, 124, 1086–1105.

OMMEREN, J. VAN, AND P. RIETVELD. (2005): "The commuting time paradox," *Journal of Urban Economics*, 58, 437–54.

OMMEREN, J. VAN, P. RIETVELD, AND P. NIJKAMP. (1997): "Commuting: In Search of Jobs and Residences," *Journal of Urban Economics*, 42, 402–21.

---. (1999): "Job Moving, Residential Moving, and Commuting: A Search Perspective," *Journal of Urban Economics*, 46, 230–53.

PERI, G., K. SHIH, AND C. SPARBER. (2015): "STEM workers, H-1B visas, and productivity in US cities," *Journal of Labor Economics*, 33, S225--S255.

Quillian, L. (2003) : "How long are exposures to poor neighborhoods? The long-term dynamics of entry and exit from poor neighborhoods", Population Research and Policy Review, 22(3), 221–249.

ROSENTHAL, S. S. (2014): "Are Private Markets and Filtering a Viable Source of Low-Income Housing? Estimates from a 'Repeat Income' Model," *American Economic Review*, 104, 687–706.

ROSS, S., AND Y. ZENOU. (2008): "Are shirking and leisure substitutable? An empirical test of efficiency wages based on urban economic theory," *Regional Science and Urban Economics*, 38, 498–517.

ROUWENDAL, J., AND P. RIETVELD. (1994): "Changes in Commuting Distances of Dutch Households," *Urban Studies*, 31, 1545–57.

SIMONSOHN, U., J. P. SIMMONS, AND L. D. NELSON. (2020): "Specification curve analysis," *Nature Human Behaviour*, .

TIMOTHY, D., AND W. C. WHEATON. (2001): "Intra-Urban Wage Variation, Employment Location, and Commuting Times," *Journal of Urban Economics*, 50, 338–66.

TSIVANIDIS, N. (2018): "The aggregate and distributional effects of urban transit infrastructure: Evidence from bogotá's transmilenio," *Mimeo*, .

WHEATON, W. C. (1977): "Income and Urban Residence: An Analysis of Consumer Demand for Location," *The American Economic Review*, 67, 620–31.

ZENOU, Y. (2009): *Urban Labor Economics*, . Cambridge BooksCambridge University Press.

Tables

Table 1 - Descriptive statistics

Panel A – Full sample

Variable	Mean	Standard Dev.	Min.	Max.	N
Worker-level variables					
Distance	18.98	15.91	1.44	175.745	1802792
Travel Time	38.80	27.80	0	402.55	1802792
Yearly Wage	24496.48	26854.59	1	7120304	1802792
Age	38.09	10.93	16	92	1802792
Gender (1 = Male; 2 = Female)	1.46	0.50	1	2	1802792
Skilled (0 = unskilled; 1 = skilled)	0.29	0.45	0	1	1802792
Mover (0 = house stayer; 1 = house mover)	0.15	0.36	0	1	1802792
Firm-level variables					
Assets	51591.18	567483.2	1	8.03e+07	30419
Apparent Labor Prod.	100.21	807.64	.11	77200.75	30419
Employment	159.0	1152.72	1	60713	30419
Total Sales (in Euros)	56801.78	391959.3	2	1.48e+07	30419
Exports (in Euros)	5199879	3.68e+07	9	1.57e+09	30419
Number HS6 products Exported	18.31	41.97	1	1090	30419

Panel B- Movers samples

Variable	Mean	Standard Dev.	Min.	Max.	N
Worker-level variables					
Distance	19.29	16.14	1.44	175.75	278211
Travel Time	39.11	27.9	0	402.55	278211
Yearly Wage	22201.1	22525.83	1	1708188	278211
Age	34.91	10.27	16	86	278211
Gender (1 = Male; 2 = Female)	1.47	0.5	1	2	278211
Skilled (0 = unskilled; 1 = skilled)	0.28	0.45	0	1	278211
Firm-level variables					
Assets	64498.89	647659.97	22	80255664	23126
Apparent Labor Prod.	96.6	742.43	0.14	77200.75	23126
Employment	201.37	1317.21	1	60713	23126
Total Sales (in Euros)	71288.65	441607.05	2	1.48e+07	23126
Exports (in Euros)	6268656.9	39757636	23	1.52e+09	23126
Number HS6 products Exported	21.01	46.65	1	1090	23126

Table 2 – Baseline IV 2SLS Results – Short-run

	(1)	(2)	(3)	(4)	(5)
First-stage results					
Dependent Var.: Log Yearly Net Wage					
Log WID (Current Weights)	0.017***	0.016***	0.016***	0.014***	0.011***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
F-stat	1243.94	1003.67	1194.48	855.63	552.55
Second-stage results					
Dependent Var.: Log Commuting Distance					
Log Yearly Net Wage	0.123***	0.163***	0.143***	0.186***	0.167***
	(0.015)	(0.015)	(0.017)	(0.020)	(0.023)
Worker FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Firm-Year Controls		✓		✓	✓
Worker Controls			✓	✓	✓
Occupation FE					✓
Observations	1802792	1802792	1802792	1802792	1802792

Notes: The WID is the world import demand for the products sold by the firm. See text for further details. Robust standard errors in parentheses with ***, **, and *respectively denoting significance at the 1%, 5% and 10% levels. Standard errors are clustered at the sector-year level. Firm-Year controls include log Assets, Log Apparent Labor Productivity, defined as total value added per worker, and log Employment. Worker controls include a skilled/unskilled dummy, the age of the worker and gender.

Table 3 – OLS baseline results – Short-run

Dependent Var.: Log Commuting Distance					
	(1)	(2)	(3)	(4)	(5)
Log Yearly Net Wage	0.026***	0.024***	0.013***	0.011***	0.006***
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)
Worker FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Firm-Year Controls		✓		✓	✓
Worker Controls			✓	✓	✓
Occupation FE					✓
Observations	1802792	1802792	1802792	1802792	1802792
R-squared	0.003	0.005	0.017	0.019	0.012

Notes: Robust standard errors in parentheses with ***, **, and *respectively denoting significance at the 1%, 5% and 10% levels. Standard errors are clustered at the sector-year level. Firm-Year controls include log Assets, Log Apparent Labor Productivity, defined as total value added per worker, and log Employment. Worker controls include a skilled/unskilled dummy, the age of the worker and gender.

Table 4 – Baseline IV 2SLS Results – Long-run

	(1)	(2)	(3)	(4)	(5)
First-stage results					
Dependent Var.: Log Yearly Net Wage					
Log WID (Current Weights)	0.039***	0.028***	0.035***	0.025***	0.016***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
F-stat	912.85	469.62	980.92	458.46	224.16
Second-stage results					
Dependent Var.: Log Commuting Distance					
Log Yearly Net Wage	0.196***	0.269***	0.215***	0.299***	0.323***
	(0.026)	(0.033)	(0.028)	(0.039)	(0.062)
Worker FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Firm-Year Controls		✓		✓	✓
Worker Controls			✓	✓	✓
Occupation FE					✓
Observations	280822	280822	280822	280822	280822

Notes: The WID is the world import demand for the product sold by the firm. See text for further details. Robust standard errors in parentheses with ***, **, and *respectively denoting significance at the 1%, 5% and 10% levels. Standard errors are clustered at the sector-year level. Firm-Year controls include log Assets, Log Apparent Labor Productivity, defined as total value added per worker, and log Employment. Worker controls include a skilled/unskilled dummy, the age of the worker and gender.

Table 5 – OLS baseline – Long-run

Dependent Var.: Log Commuting Distance	(1)	(2)	(3)	(4)	(5)
Log Yearly Net Wage	0.067***	0.061***	0.038***	0.033***	0.013***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Worker FE	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓
Firm-Year Controls		✓		✓	✓
Worker Controls			✓	✓	✓
Occupation FE					✓
Observations	285106	285106	285106	285106	285106
R-squared	0.008	0.009	0.019	0.020	0.010

Notes: Robust standard errors in parentheses with ***, **, and *respectively denoting significance at the 1%, 5% and 10% levels. Standard errors are clustered at the sector-year level. Firm-Year controls include log Assets, Log Apparent Labor Productivity, defined as total value added per worker, and log Employment. Worker controls include a skilled/unskilled dummy, the age of the worker and gender.

Table 6 – Location Choice

Dependent Variable:	Distance to Paris	Average Revenue 2002	Workers Density 2002	Firm Density 2002	Exporters Density 2002	Travel Time Ratio
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: short-run estimates						
Log Yearly Net Wage	-0.015***	0.015***	0.081***	-0.020***	-0.025***	0.053***
	(0.005)	(0.004)	(0.004)	(0.005)	(0.005)	(0.010)
Quantification: … as % of st.d of dependent variable	*17,85%*	*59,5%*	*45%*	*18%*	*25%*	*11,5%*
Worker FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Firm-Year Controls	✓	✓	✓	✓	✓	✓
Worker Controls	✓	✓	✓	✓	✓	✓
Observations	1597288	1597288	1597288	1597288	1597288	1597288
Panel B: long-run estimates						
Log Yearly Net Wage	-0.013***	0.025***	0.101***	-0.029***	-0.019***	0.084***
	(0.004)	(0.008)	(0.020)	(0.006)	(0.006)	(0.023)
Quantification: … as % of st.d of dependent variable	*15,29%*	*61%*	*55%*	*30,2%*	*18,5%*	*16,7%*
Worker FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓	✓	✓	✓
Firm-Year Controls	✓	✓	✓	✓	✓	✓
Worker Controls	✓	✓	✓	✓	✓	✓
Occupation FE						
Observations	276502	276502	276502	276502	276502	

Notes: Second-stage IV estimations, in which World Import Demand and World Demand measures (see text) are used as excluded instruments for yearly net wages. Robust standard errors in parentheses with ***, **, and *respectively denoting significance at the 1%, 5% and 10% levels. Standard errors are clustered at the sector-year level. Firm-Year controls include log Assets, Log Apparent Labor Productivity, defined as total value added per worker, and log Employment. Worker controls include a skilled/unskilled dummy, the age of the worker and gender.

Table 7 – Omitted Variables and Additional Fixed Effects

Short-run Estimates

Dependent Var.: Log Commuting Distance						
	(1)	(2)	(3)	(4)	(5)	(6)
Log Yearly Net Wage	0.264***	0.272***	0.348***	0.139***	0.155***	0.319***
	(0.020)	(0.016)	(0.019)	(0.029)	(0.056)	(0.013)
Firm-Year Controls	✓	✓	✓	✓	✓	✓
Worker Controls	✓	✓	✓	✓	✓	✓
Worker FE	✓	✓	✓	✓	✓	✓
Year FE	✓					
House City FE	✓					
House City x Year FE		✓		✓	✓	✓
House City x Year x Sector FE			✓			
Employment Zone FE				✓		
Employment Zone x Year FE					✓	
4-digit Industry x Year trends						✓
Observations	1802792	1802792	1802792	1802792	1802792	1802792

Long-run Estimates

Dependent Var.: Log Commuting Distance						
	(1)	(2)	(3)	(4)	(5)	(6)
Log Yearly Net Wage	0.412****	0.426***	0.546***	0.229***	0.264***	0.422***
	(0.025)	(0.029)	(0.019)	(0.029)	(0.039)	(0.017)
Firm-Year Controls	✓	✓	✓	✓	✓	✓
Worker Controls	✓	✓	✓	✓	✓	✓
Worker FE	✓	✓	✓	✓	✓	✓
Year FE	✓					
House City FE	✓					
House City x Year FE		✓		✓	✓	✓
House City x Year x Sector FE			✓			
Employment Zone FE				✓		
Employment Zone x Year FE					✓	
4-digit Industry x Year trends						✓
Observations	280822	280822	280822	280822	280822	280822

Notes: Second-stage IV estimations, in which World Import Demand (see text) is used as excluded instrument for yearly net wages. Robust standard errors in parentheses with ***, **, and *respectively denoting significance at the 1%, 5% and 10% levels. Standard errors are clustered at the sector-year level. Firm-Year controls include log Assets, Log Apparent Labor Productivity, defined as total value added per worker, and log Employment. Worker controls include a skilled/unskilled dummy, the age of the worker and gender.

Figures

Figure 1 - Specification curve (Simonsohn et al. (2020)) - Short-run estimates

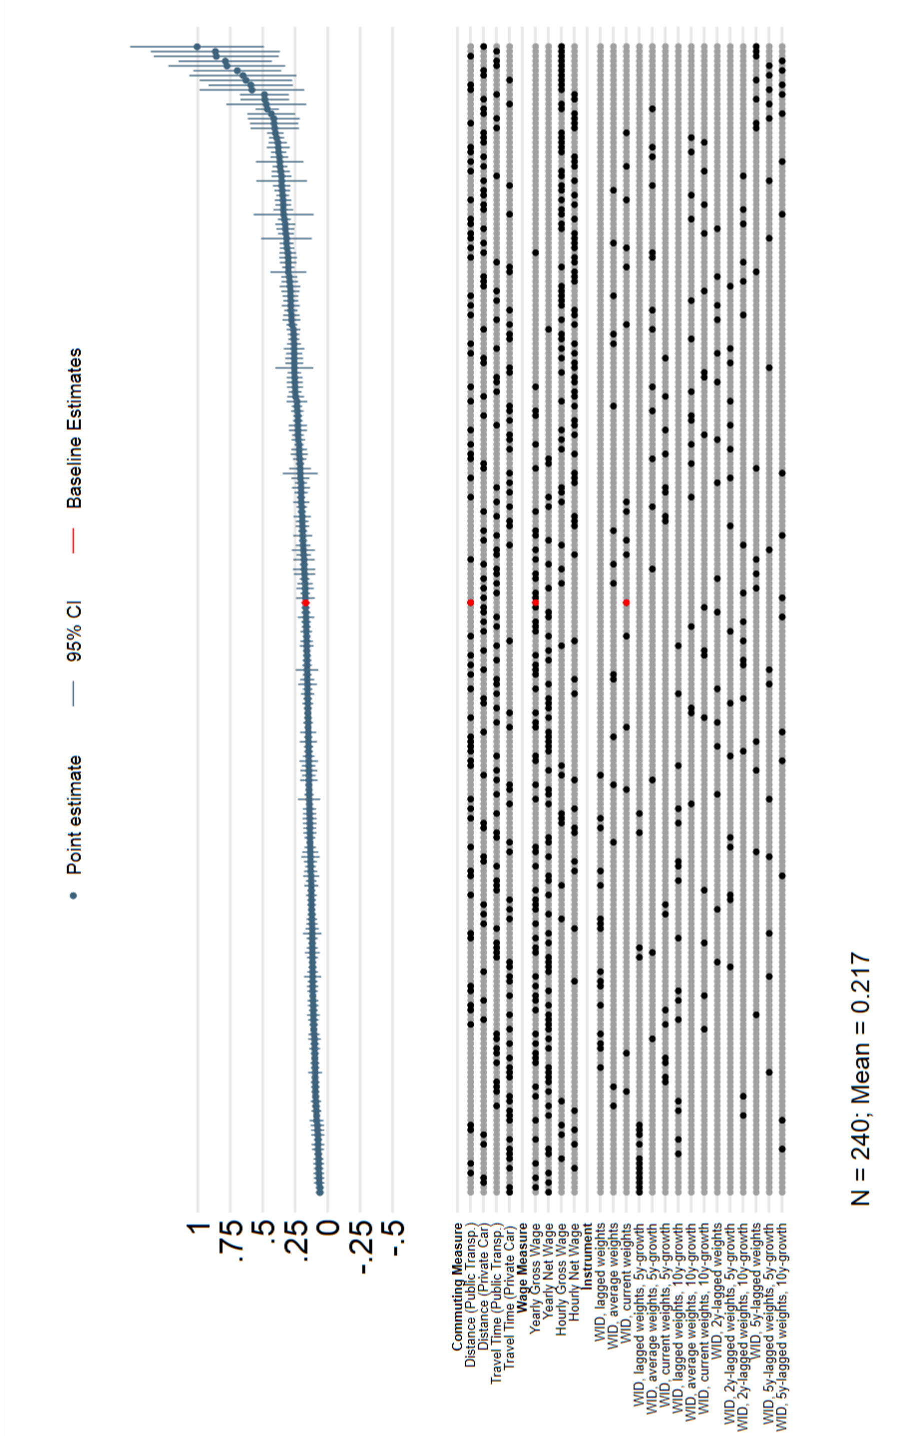

Figure 2 - Specification curve (Simonsohn et al. (2020)) – Long-run estimates

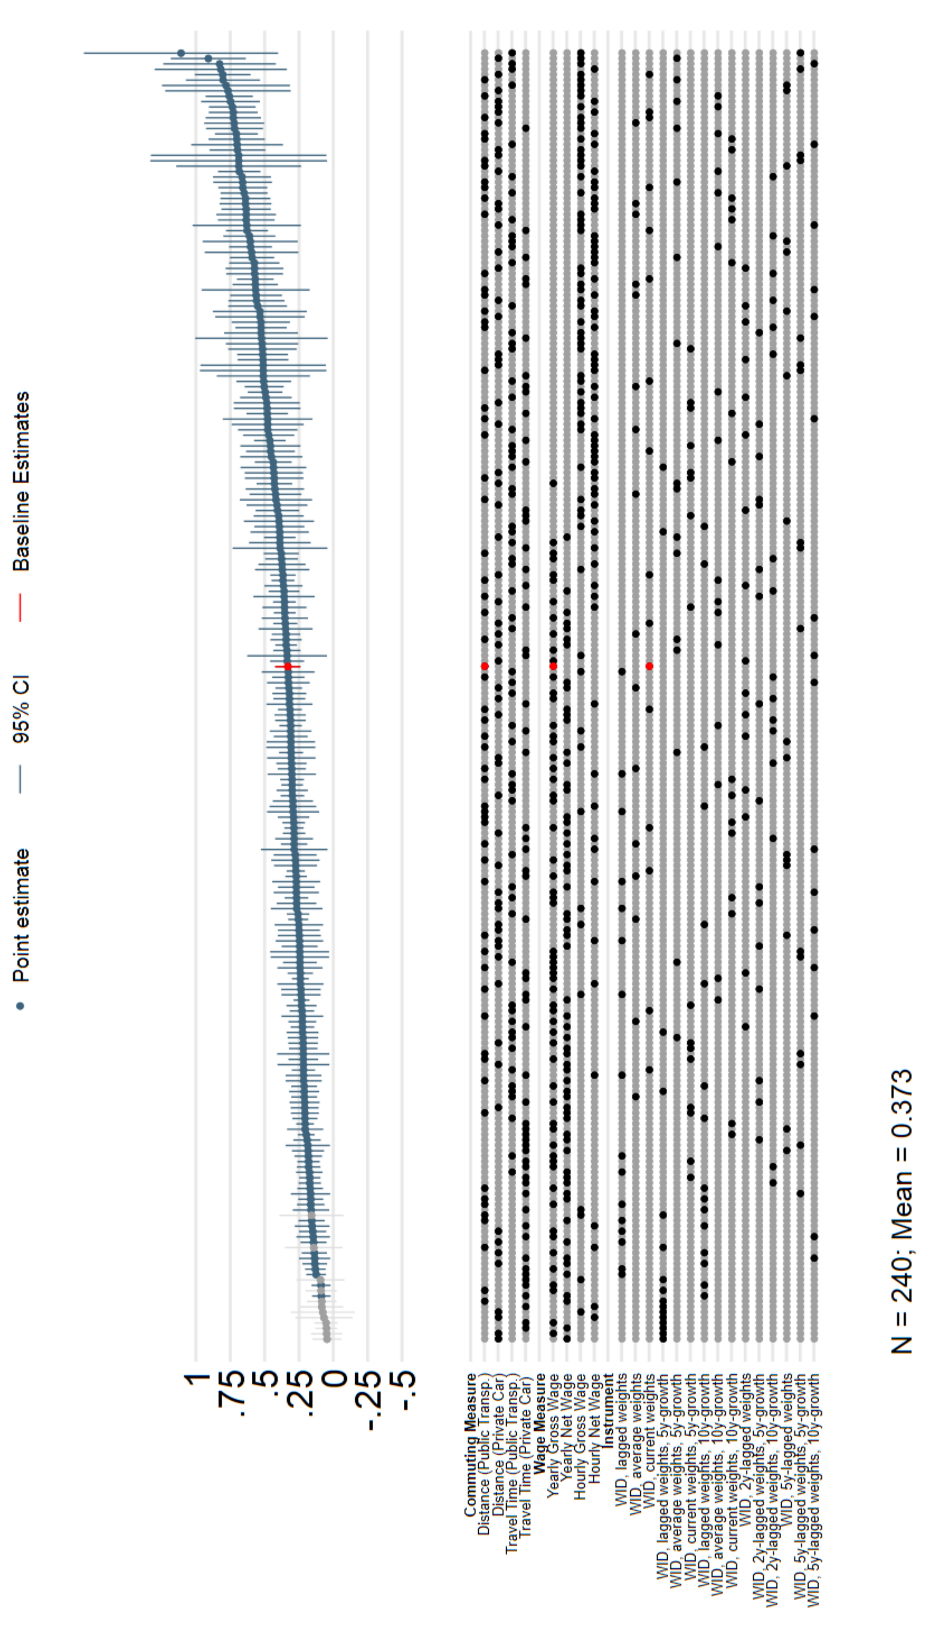

Figure 3: Event study design