

Is Tourism good for Locals? Evidence from Barcelona *

Treb Allen

Dartmouth and NBER

Sharat Ganapati

Georgetown University

Rocio Madera

Southern Methodist Univ.

Simon Fuchs

Atlanta Federal Reserve

Alberto Graziano

CaixaBank Research

Judit Montoriol-Garriga

CaixaBank Research

July 2020

Abstract

Is tourism good for locals? We embed a Ricardo-Viner framework into a rich urban geography and show that the welfare impact of shocks depends only on (1) the spatial patterns of consumption and income; and (2) the price and wage effects of the shock throughout the city. We use spatially disaggregated consumption and income data to estimate the price and wage effects of Barcelona's tourist boom. We identify these effects using an identification strategy based on monthly variation in the aggregate composition of tourists' origin. We find that, on average, local workers suffer slightly from tourism, but these average effects mask substantial heterogeneity across space, ranging from a -19 to a +4 percent welfare change between low and high tourist seasons. The inner city residents bear the largest price changes but enjoy substantial income gains, whereas peripheric neighborhoods suffer lower but sizable price changes with none to moderate income benefits.

*Preliminary and incomplete. We are grateful to Cecile Gaubert for her excellent discussion and to Javier Ibañez de Aldecoa Fuster for his excellent support with the data. The views expressed herein are those of the authors and not necessarily those of CaixaBank, the Federal Reserve Bank of Atlanta, or the Federal Reserve System.

1 Introduction

In many locations around the world, tourism comprises a substantial and growing fraction of the local economy. For example, tourism is Spain’s single largest export sector and the second fastest growing sector of the economy, with tourism expenditures currently equal in value to half of all Spain’s exports of goods (and 11% of GDP in total). The rise in tourism appears to be driven in large part by increased international demand and falling travel costs; for example, the advent of low cost airlines has resulted in the air passenger volume within the E.U. to have tripled over the past 25 years.

In a standard trade model, an increase in foreign demand for a location’s export good should be welfare improving for the residents of that location. Yet in many cities, tourism is immensely unpopular among local residents (see e.g. Figure 1), who complain that tourists “crowd out” locals, raising the prices of local consumables. In Barcelona, for example, frequent protests against tourism have occasionally led to altercations between locals and tourists, and the current mayor campaigned on a promise to limit tourism in the city. Are standard trade models wrong? Or are residents complaints misplaced?

In this paper, we answer the question “Is tourism good for locals?” To do so, we develop an urban Ricardo-Viner framework featuring a rich geography of consumption and income patterns and show that the welfare impact of any shock depends on (1) the spatial patterns of consumption and income of locals throughout the city; and (2) how all prices and wages throughout the city change in response to the shock. We assemble a new high-resolution spatial dataset on consumption and income patterns in the city of Barcelona which allows us to observe the former. We then combine a novel identification strategy relying on aggregate variation in the composition of the country of origin of tourists and general equilibrium market clearing conditions to estimate the latter.

On average, we find that tourism is slightly negative for locals; however, these (modest) average losses mask substantial heterogeneity all across the city: The welfare changes from moving between low and high tourist seasons range from a *negative* 19 percent (10th percentile in the welfare changes’ distribution over locations) to a *positive* 4 percent (percentile 90th). Dissecting this net welfare changes into price and income effects, residents in the city center and those near tourist locations bear the largest price changes but also enjoy substantial income gains. In contrast, residents of peripheric neighborhoods suffer lower but still sizable price changes, with the income gains varying between different outer city locations: some experience none and some get moderate income benefits from tourism.

The data we assemble is based on hundreds of millions of credit and debit card transactions recorded by one of the largest banks in Spain and dominant bank in Barcelona between

January 2017 and December 2019, covering roughly 3% of the entire Barcelona metro area GDP. The transactions are from two sources: (1) purchases made at a point-of-sale owned by the bank; and (2) purchases made by customers of the bank. The former allows us to construct monthly level expenditure for 20 different product categories across 1,095 locations within the Barcelona metro area for tourists from each non-Spanish country. The latter allows us to construct *bilateral* expenditure share matrices for Barcelona residents by both their location of residence and the location and category of purchase. (It also allows us to construct expenditure data for non-local Spanish tourists). To account for the fact that not all purchases are made using a credit or debit card, we append additional housing rental data and re-weight the expenditures by product category to match aggregate expenditure surveys. We combine this detailed expenditure data with commuting data at the same spatial resolution, which we construct by cross-referencing from cell-phone location data with commuting survey data.

Using this dataset, we first document three stylized facts. First, tourist expenditure varies substantially across the city and over time. While certain locations are always popular amongst tourists, the relative popularity of locations depends importantly on the country of origin of the tourist (e.g. Spanish tourists prefer shopping malls, whereas international tourists prefer the beaches). This fact will be particularly helpful in the empirics, as it allows us to use aggregate variation in the composition of tourists in the city at a given time – driven e.g. by differences in timings of school breaks in the origin countries – to generate variation in tourism expenditure that is plausibly orthogonal to unobserved changes in local conditions. Second, we document that both local expenditure and income has a strong “gravity” spatial component, i.e. locals are much more likely to purchase goods and work nearby their residence. Combined with the first stylized fact, this implies that residents living closer to places popular with tourists will be more exposed to tourism. Third, comparing the tourism “low” and “high” seasons within a year, we show that total sales increase more in locations popular with tourists where tourism expenditure increases but that local expenditure falls the most in these same locations, suggesting that tourism both increases incomes earned by locals but also increases the prices locals pay for goods.

To understand the welfare implications of these price and income changes, we embed a Ricardo-Viner specific factors trade model into an urban setting with a rich geographic patterns of consumption and commuting. We model tourism as an increase in foreign demand for goods produced only in Barcelona, where goods are differentiated based on both on their type and where in Barcelona they are produced (allowing us to capture different tourism preferences for different tourist sights within the city). We consider three variants of our framework. In the first “simple” variant with only two sectors (a tourist sector and a tradable

sector), we show that locals benefit from a small increase in tourism if and only if they are net producers in the tourist sector, in which case the welfare impact of the income gains exceed the welfare costs of price increases.

This basic intuition extends readily to the second “general” variant where we introduce a rich geography of consumption and income: as in the simple model, the welfare impact of tourism depends on whether it increases a residents’ income more than it increases the prices she pays for her goods. However, these two effects now depend on both that particular resident’s patterns of consumption and earnings and, crucially, on how wages and prices change in response to tourism throughout the city. While spatial patterns of consumption and earnings are observed in the data, these wage and price effects are not. However, by imposing standard market clearing conditions, we derive simple analytical expressions for the (short-run) wage and price effects that depend only on observed data.

To determine the (long-run) prices and wage effects from potentially large changes in tourism expenditure, in the third “quantitative” variant of our model, we impose functional form assumptions on production, consumption, and commuting. As in the previous variant, the welfare impact of tourism depends on the observed spatial patterns of consumption and earnings and how prices and wages change throughout the city. The price and wage effects can then be calculated from market clearing conditions given observed spatial patterns and knowledge of supply and demand elasticities.

Combining the data with our theoretical frameworks, we then empirically evaluate the welfare impact of tourism on locals. We pursue three distinct but complementary evaluation strategies, which vary in the relative weights they place on theory and data to recover how local wages and prices respond to increases in tourism. In the first “deductive” approach, we rely entirely on data to estimate the average wage and price effects for the city as a whole. To do so, we employ a shift-share methodology that takes advantage of the aggregate variation over time in the composition of tourists from different countries of origin to generate plausibly exogenous variation across the city in tourism expenditure. Consistent with the stylized facts above, we document that tourism casually increases local wages and increases local prices. Given observed expenditure shares, the positive impact on prices dominates the impact on wages wages, i.e. we find that the average resident in Barcelona suffers slightly from tourism.

However, by relying on average wage and price effects for the city as a whole, the deductive approach potentially masks substantial heterogeneity in who wins and loses from tourism across the city. In the second “inductive” approach, we instead rely on the derivations from our theoretical framework to recover the wage and price effects throughout the city. In the both the “short-run” (i.e. holding labor allocations and expenditure shares fixed) and the

“long-run” (where we allow labor and expenditure shares to respond to tourism shocks), we find substantial heterogeneity in the welfare impacts of residents across space. Residents of the center of Barcelona (where much of the tourism occurs) are hurt by tourism, as the price increases dominate their wage gains. In contrast, residents of the outskirts of Barcelona are protected from the negative effects of tourism, suffering only modest negative price effects (as they mostly consume in non-tourist areas) and positive wage effects (as many commute to locations where wages have risen). Their net welfare effect tends to be close to zero.

The disadvantage of the inductive approach is that it requires imposing additional assumptions: to estimate wage and price effects in the short run, we must impose a market clearing condition, and in the long run we also rely particular functional form assumptions for production and consumption responses. In our third “deductive-inductive” strategy, we combine the previous two approaches. We interact our theory-predicted wage from the inductive approach in the price and wage regressions from the inductive approach to evaluate the extent to which the theory is able to predict the variation price and wage effects across the city in the data. We find that the theory does a good job of capturing the variation across space in price and wage effects. Combined, all three empirical approaches paint a robust and consistent pattern of the welfare impacts of the tourism on locals, with overall modest gains that are unequally shared among residents across space.

This paper makes three primary contributions to the literature. First, we provide an estimate of the spatially heterogeneous welfare impact of tourism on locals throughout a city. While several recent papers have examined the impact of tourism on local housing markets and consumption amenities (e.g. [Almagro and Domínguez-Iino \(2019\)](#) and [García-López, Jofre-Monseny, Mazza, and Segú \(2019\)](#)), they have tended to abstract from spatial linkages (through either commuting or consumption) within the city, instead treating different neighborhoods as independent locations. Here, we explicitly model these linkages and show they play an important role generating heterogeneity in wage and price effects across the city. In this way, the paper is closely related to [Faber and Gaubert \(2019\)](#), who show that the welfare impact of tourism depends importantly on spatial and sectoral linkages, albeit across regions within a country instead of neighborhoods within a city.

Second, we extend the welfare results of a specific factors Ricardo-Viner trade model (see e.g. [Mussa \(1974\)](#); [Jones \(1975\)](#); [Mussa \(1982\)](#) for analysis of a single location-country and [Kovak \(2013\)](#); [Dix-Carneiro and Kovak \(2017\)](#) for analysis of multiple region-countries) to urban settings with rich geographies where agents move across space to both consume and produce. Despite complex spatial consumption and commuting patterns, it turns out that all one needs to calculate the welfare impact of any economic shock is: (1) the spatial patterns of consumption and income of residents; and (2) how the shock affects prices and wages

throughout the city. As large-scale spatial data sets become increasingly available, (see e.g. [Athey, Blei, Donnelly, Ruiz, and Schmidt \(2018\)](#) and [Couture, Dingel, Green, and Handbury \(2020\)](#) for examples using mobile phone data, [Davis, Dingel, Monras, and Morales \(2019\)](#) for an example using online review data, and [Carvalho, R. Garcia, Hansen, Ortiz, Rodrigo, Rodriguez Mora, and Ruiz \(2020\)](#) and [Agarwal, Jensen, and Monte \(2017\)](#) examples using credit card transaction level data), we expect the first ingredient will become increasingly attainable. To assist in calculating the second ingredient, we also provide new and intuitive analytical expressions for how prices and wages respond to a shock in the short-run (relying on tools introduced by [Allen, Arkolakis, and Takahashi \(2020\)](#) for trade models) and in the long-run (relying on tools introduced by [Dekle, Eaton, and Kortum \(2008\)](#) and detailed in [Costinot and Rodriguez-Clare \(2013\)](#) for trade models).

The third contribution of the paper is to propose a new empirical methodology that marries recent advances in the *quantitative* spatial literature with recent advances in the *applied* spatial literature. While the seminal paper of [Ahlfeldt, Redding, Sturm, and Wolf \(2015\)](#) introduced general equilibrium counterfactual analysis to urban models with complex geographies, retaining tractability required making particular functional form assumptions on preferences, commuting, and production. Here, we relax those assumptions in two ways: first, we show that in the “short-run” and for small shocks, price and wage changes can be calculated with no functional form assumptions (beyond homothetic demand and constant returns to scale production); second, in the spirit of [Donaldson \(2018\)](#) and [Monte, Redding, and Rossi-Hansberg \(2018\)](#), we analyze the extent to which the model predicted price and wage effects are able to capture the observed empirical variation in prices and wages.

The rest of the paper is organized as follows. The next section describes our data and presents three stylized facts. Section 3 presents three variants of our theoretical framework to show theoretically how tourism affects locals’ welfare. Section 4 combines the data and theory in three approaches to empirically evaluate the welfare impacts of tourism. Section 5 concludes.

2 Tourism in Barcelona

Tourism is a key sector in Catalonia, and Barcelona in particular. In 2018, 19.12M Foreign tourists visited Catalonia (Idescat 2019), approximately doubling in a decade. On average each tourist spent €185 per day, totaling in €20.6B in declared expenditures. It is safe to assume that the largest share of these expenditures were spent in the tourist hotspots of Barcelona, but detailed data on tourist economic activities at the microgeographic level is difficult to come by, with commonly available datasets only giving a coarse understanding of

the impact of tourism.

For this project we draw on multiple data sources that describe in fine geographical detail economic activities of tourists and locals within the city of Barcelona. The most important data source is a new expenditure database that is constructed from electronic payments processed by CaixaBank. In this section we introduce this database that describes in much detail expenditures in Barcelona, both by tourists and locals. We will use this novel data source to derive some key stylized facts that describe tourism as an urban spatio-temporal phenomenon. These stylized facts will motivate our theoretical framework and our empirical strategy. We will also introduce a set of additional additional data source that we will use in our structural exercise. This includes multiple data sources describing commuting flows within the city of Barcelona: (1) A traditional commuter survey, (2) lunch-time expenditures on weekdays derived from our expenditure database and (3) commuting flows from cell-phone locations as processed by INE. We will also use detailed information on time trends of rental rates at the neighborhood level obtained from Idealista (a local real estate aggregator).

2.1 New high resolution spatial and temporal data on local and tourist expenditure

Expenditure data

We use transaction-level data from electronic payments that were submitted to Caixabank's Payment Processing Service. Caixabank is the leading bank among individuals and SMEs in Spain and is based in Barcelona, where it has close to a 40% market share. The underlying data contains each debit or credit card purchase at any merchant with a Caixabank Point of Sale (PoS) in the city of Barcelona. For each transaction, the total euro amount, the exact merchant geo-localization, the expenditure category, the country of origin of the paying credit card, as well as the time and date when it happened, are recorded. Importantly for us, if the customer is a Caixabank client herself, her home address is additionally registered, allowing us to trace out the spatial expenditure pattern of a residential location in Barcelona. Our data of analysis consist of the value of the full set of these transactions per month and census-tract in the Barcelona metropolitan area (Àmbit metropolità de Barcelona), further disaggregated by merchant category, type of customer (resident or tourist, and subgroups within), and origin location of the customer (census-tract if resident and country if tourist).

To put the scope of our data in context, we have over 165 million yearly observations adding up to a total value of 2,970 million euros. There are 1068 census tracts in the city of Barcelona proper and we further include the 27 municipalities that form the metropolitan area of Barcelona (AMB). Our data span from January 2017 through December 2019.

We define five customer groups based on residence status and data availability (Caixabank relation): (1) residents that are Caixabank customers, (2) residents that are not Caixabank customers, (3) domestic tourists that are Caixabank customers, (4) domestic tourists that are not Caixabank customers, (5) and foreign tourists. The latter group can be further disaggregated by country of origin based on the credit card bank, but we treat it as a unique group for now. We document detailed destination-level descriptive analytics for all groups. In our benchmark analysis, however, we define resident transactions as those originating from members of group (1), and tourist transactions as those originating from members of groups (3) through (5). We might further distinguish in the analysis between domestic tourists comprising groups (3) and (4), and foreign tourists, group (5).

To ease concerns about the representativeness of expenditures taken from electronically processed transactions only (as opposed to cash transactions), we have created a crosswalk from the original coding in terms of merchant categories to COICOP - a classification of expenditures into consumption categories commonly used for expenditure surveys. This allows us to directly compare the expenditure shares in our sample to the national expenditure survey in Spain as is down in table 4. Our dataset corresponds to 54.4pc of the expenditures observed in national expenditure survey. The weights on individual categories is roughly comparable to the results of the expenditure survey, but far from exactly matching it. In our structural general equilibrium analysis we apply a re-weighting strategy of the residential expenditure patterns to make it coherent with the expenditure survey and representative for local expenditure pattern overall. For tourist expenditure pattern no such adjustment was applied.

Housing Rental rate data

While detailed in many aspects, our expenditure database is not informative on housing expenditure, a category of expenditures not commonly payed using debit or credit cards in Spain. We therefore employ an additional database on local rental rates across Barcelona that we obtained from Idealista. Idealista imputes rental rate trends at a monthly frequency for neighborhoods (Barrios) in Barcelona. Neighborhoods contain multiple census sections and therefore the geographical resolution of this data is coarser than our expenditure data, but nevertheless it is sufficient to capture key trends and cross-sectional variation in the Barcelona housing market. The sample for the housing data covers the period between January 2010 and June 2020.

Commuting data

Commuting data has traditionally been a cornerstone of applied urban analysis. Traditionally, surveys have been used to inform our understanding of employment linkages across the urban landscape. Our first dataset describing commuting in Barcelona is such a survey called L'Enquesta de Mobilitat en Dia Feiner (EMEF), and is conducted annually by the municipality of Barcelona, the AMB, the transport authorities (ATM) and the association for mobility and urban transport (AMTU).

More recently, additional data on urban mobility has become available exploiting the spatial extent of the cell phone network. As cell phones and their owners chart their path through the city, they continue to switch between cell phone towers, logging into the closest one to optimize network coverage. This leaves a data trail of time stamps and associated towers for each cell phone which in turn can be used to impute the spatial path of cell phones through the day. In the light of the Covid-19 crisis, the national statistical agency in Spain has acquired data from the most important network operators in Spain to impute the mobility patterns of residents across Spain. They also released a benchmark dataset that describes mobility patterns on November 18th, 2019. We use this data to inform our analysis. The data describes the flows between spatial units. The geographical aggregation of this data is somewhat coarser, being coded at the neighborhood (barrio) level. An additional challenge is that for privacy reasons INE does not report bilateral flows that are less than 100 in absolute magnitude.

Finally, we also constructed our own commuting flows from the expenditure database above. The general idea is to analyze the subset of the data where we observe the residential location of the account holder and to isolate their lunchtime expenditures on weekdays. Assuming that lunchtime expenditures are very proximate to the place of work, this strategy can possibly isolate commuting flows. An interesting advantage of this approach is that this can be done at the same geographic resolution as the expenditure data, i.e. it allows us to recover commuting patterns between census blocks.

In our current analysis we mostly rely on the cell-phone location derived data, judging it to be of the highest quality amongst the available dataset. Commuting flows from lunchtime expenditures remain too noisy to be directly used, but we hope that future work relying on a structural gravity imputation approach can combine the reliability of the cellphone location data with the higher geographical resolution of the expenditure imputed data.

2.2 Three stylized facts

In this section we use our data sources to document three stylized facts: (1) tourism varies across space and time within the city; (2) local's consumption and income exhibits strong spatial patterns localized around their place of residence; and (3) tourism appears to crowd out local consumption (but increase total spending), consistent with it having both price and wage effects.

Fact #1(a): Tourism is spatial: it varies substantially across space (and across sector)...

Figure 2 shows the intensity of tourist expenditures across individual locations for a given year. The intensity is normalized by the area of the underlying tile to account for heterogeneity in the size of individual census blocks. For convenience, we also show - with blue labels - the location of 15 of the most popular tourist sites in Barcelona. Not surprisingly the expenditure of tourists is closely correlated with the location of the main tourist attraction. The historical medieval core of Barcelona together with its extension towards Gracia forms an axis of intense touristic activity, with additional hotspots close to La Sagrada Familia and along the beachfront. Notice, that our data is sensitive to the intensity of commercial activity across different locations. We indicate with yellow labels the largest shopping centers that tend to be associated with high levels both for tourists and - as we will show further below - for residents. Overall, however it is clear that tourist activity is fairly concentrated in the historical core of the city.

Fact #1(b): Tourism is spatial: ... but the spatial incidence varies over time within the year (depending on origin of tourists)

In Figure 3, we map as a bivariate choropleth the quantiles of Spanish and foreign tourists across Barcelona. Spanish tourists are defined as Spanish account holders that visit from outside of the province of Catalonia. In the map locations that experience higher foreign tourist expenditure, but comparatively lower domestic tourist expenditures are marked with green colors, while the reverse situation - low foreign but high domestic tourist expenditures - are marked in magenta tones. Locations that experience both high expenditures by domestic and foreign tourists are marked in dark grey, while low expenditure locations are marked in light gray. The inner city is popular with both domestic and foreign tourists. While Baraceloneta and Montjuic are particularly popular with foreign visitors, some of the outer areas tend to be more popular with domestic visitors. Overall, there is evidence for a distinctive heterogeneity in preferences for locations between different groups of tourists.

The spatial heterogeneity of different tourist groups interacts with their importance across time, in particular their seasonality. In Figure 4 and in table3 we demonstrate the expenditure composition in our data. What is probably most striking is the seasonal variation in total tourist expenditures. Between February and July on average, combined expenditures of domestic and foreign tourists increase by a stunning 70 percent. Foreign tourist expenditures are particularly volatile with a clear seasonal pattern with a distinct high season during the summer months, while domestic tourists are more stable throughout the year. This variation will prove enormously helpful when we seek to estimate price and wage effects in Section 4 below.

Fact #2(a): Consumption is spatial: Residents are more likely to spend nearby their home

Residential consumption experiences spatial decay, with the own location often being the destination of a clear majority of expenditure of residents and with other expenditures strongly declining as distance and travel cost increases. In Figure 5 we plot the expenditure shares for a resident of the historical urban core. There are two clear take aways from the figure: The first is that there is a substantial share of expenditures in close vicinity to the residential location. A large fraction of the expenditures take place in less than 1km distance from the home location. The second observation is that expenditure patterns are widely spread across the city, reaching into almost all areas.

To further emphasize the spatial pattern of consumption for residents we project the observed expenditure shares of local residents living in location n on purchases in sector s and location i , π_{nis} , on bilateral travel time τ_{ni} and sector specific origin and destination fixed effects, i.e.:

$$\log \pi_{nis} = \phi_s \log \tau_{ni} + \log \delta_{n,s} + \log \delta_{i,s} + u_{ni,s},$$

where we obtain travel time using the HERE API, averaging times between public transit and driving times. For the regression we use an average between the two measures. The resulting coefficient for distance is visualized in Figure 6. There is stark heterogeneity across sectors, emphasizing the importance of sectoral data to understand the spatial component of urban consumption.

Fact #2(b): Income is spatial: Residents are more likely to earn nearby their home

As is more commonly known, and as has previously been explored in the urban economics literature, transport cost are prohibitive and prevent long commutes, inducing people to choose locations close to their workspace, resulting in a localized employment pattern. This is similarly apparent in our data. In Figure 7, we show compare commuting patterns from the INE cell phone location data for residents near the city center and those outside the city center. Even though the data is coarser and truncated - as has been mentioned above - what is striking is the localized pattern of the commute, with the majority of the commuting trips leading to neighboring locations in the city. This suggests a very strong distance coefficient in the gravity regression for commuting.

Fact #3(a): Tourism seems to affect locals: It crowds out local consumption...

In Figure 8, we map the bivariate choropleth simultaneously showing the change in tourist and local expenditure between February (the tourist low season) and August (the tourist high season) in 2019. It is clear from the map that particularly in the (foreign) tourist hotspots close to the beach as well as the lower part of the historical inner city tourist expenditures increase while local expenditures decrease or only grow weakly. We interpret this to be suggestive evidence that tourism increased prices of local goods, crowding out local consumption. Instead residential expenditure grows much stronger in the predominantly non-tourist residential locations in the northern part of the city. The map also points towards the identification problem that underlies the challenge of estimating the crowding out effect. While some areas seem to indicate crowding out, there are also central areas where expenditures strongly co-move, indicating that - despite possible price increases effects - these areas become more attractive for both tourists and locals in certain seasons.

Fact #3(b): Tourism seems to affect locals: ... but it increases total sales

In Figure 9 we compare total sales between February (the tourist low season) and August (the tourist high season) in 2019. As is evident, despite the suggestive evidence of tourist expenditure crowding out local expenditure from Fact #3(a), total expenditure grew more rapidly in the tourist high season in heavily toured areas. Together with Fact #3(a), this is consistent with tourism both increasing local prices (and crowding out local spending) and increasing local incomes (by increasing total spending).

We now turn to developing a theoretical framework to assess the welfare impacts of tourism. Given the evidence of substantial consumption and income heterogeneity of resi-

dents across space from Fact #2, our framework incorporates a flexible urban geography in order to incorporate the complex observed patterns of local consumption and expenditure. Given the suggestive evidence of both price and income effects from Fact #3, our framework allows for tourism to have arbitrary impacts on prices and wages throughout the city.

3 Is tourism good for locals? Three theoretical frameworks

In this section, we develop a three related theoretical frameworks to answer the question “is tourism good for locals?” In each framework, we model tourism as demand by non-residents for the output of a sector (or sectors) which requires at least some labor by residents of that city as an input in its production, which allows us to characterize the welfare impacts of tourism on locals through the lens of a specific factors (Ricardo-Viner) trade model.

The section is organized as follows. We first present a simple theoretical framework to provide intuition for the economic mechanisms at play. We then present a general framework that incorporates a rich urban geography of consumption and production and use it to derive analytical expressions for the (short-run, local) welfare impacts of tourism that depends only on observed spatial patterns of consumption and production and holds for any (homothetic) preferences, any (constant returns to scale) production functions, and for general labor allocation choices. Finally, we present a quantitative framework which, by specifying a particular preference structure, production functions and commuting decision process, allows us to derive welfare impacts of tourism more generally.

3.1 A simple theoretical framework

To provide some intuition for the results that follow, we begin with a very simple model. Consider a city i (Barcelona) that is inhabited by a representative local resident endowed with R_i units of labor and is small relative to the rest of the world. Suppose there are just two sectors s in this economy: a non-tourist sector – say, manufacturing – indexed by $s = 0$ and a tourist sector – say, local restaurants – indexed by $s = 1$. We suppose the output in the non-tourist sector is costlessly traded and its price (which we treat as the numeraire) is set by the world market (i.e. $p_0 = 1$). Production in either sector requires the labor of the resident L_{is} and a specific factor M_{is} , with some constant returns to scale production function $Q_{is} = F_{is}(L_{is}, M_{is})$.

A brief discussion is necessary on our interpretation of the specific factor. In the theory, we assume the specific factor is owned by agents residing outside of Barcelona, and accordingly

focus on the welfare impact of tourism on local labor. This assumption allows the framework to incorporate the empirical fact that not all proceeds from tourism accrues to locals: for example, many of the hotels within Barcelona are earned by international firms (who are compensated for providing the necessary capital). Of course, not all capital provided is from non-residents (e.g. some restaurants are owned by residents), so in the empirics below we will be choosing the labor share to match the observed relationship between sales and local resident's expenditure, thereby incorporating non-wage proceeds to locals into the measure of returns to labor as well.

The representative resident makes two choices: first, she chooses how to allocate her labor across sectors in order to maximize her income. Second, given income and prices, she chooses her consumption to maximize their utility. We assume she has homothetic demand that can be represented in the following indirect utility function $u_i = \frac{v_i}{G(\{p_{is}\}_{s \in \{0,1\}})}$ for some price aggregator $G(\cdot)$, where v_i is the income of the resident.

We now ask the following: how does a change in tourism affect the welfare of the representative local? For now, we assume that tourism can be represented by an increase in the price of the tourist sector. To analyze the impact of such a price change, we first fully (log) differentiate the indirect utility function and apply Shepherd's lemma which yields:

$$d \ln u_i = \partial \ln v_i - \sum_s \pi_{is} \partial \ln p_{is}, \quad (1)$$

where $\pi_{is} \equiv \frac{x_{is}}{v_i}$ is the local expenditure share on sector s . Equation (1) is intuitive: it says for any shock, the first order welfare impact on locals depends on two things: first, how that shock affects local's income; and second, how that shock affects local's price of consumption, which is simply a expenditure share weighted change in prices.

Because locals allocate their labor across sectors to maximize their income, the change in income and prices is closely related. Applying the envelope theorem to workers' labor allocation decision¹ yields that the change in local income is log proportional to the change in prices, where the proportion is simply equal to locals income share from sector s , i.e.:

$$\frac{\partial \ln v_i}{\partial \ln p_{is}} = \sigma_{is},$$

¹Here we assume the representative agent accounts for her labor allocation's effect on her marginal product of labor and, hence, her wage (rather than taking her wage as given). This simplifies the following discussion, although at a cost of verisimilitude; as a result, we will not make this assumption in the general empirical framework we employ empirically (and which we present next).

where $\sigma_{is} \equiv \frac{w_{is}L_{is}}{v_i}$. As a result, we can express the welfare impact on locals as follows:

$$d \ln u_i = \sum_s (\sigma_{is} - \pi_{is}) \partial \ln p_{is} \quad (2)$$

Equation (2) is also intuitive: it says locals benefits from a price increase if and only if they are net producers of the good.

In this simple example with only two sectors and the fact that wages are equalized across sectors allows us to re-write the expression even more simply:

$$\frac{\partial \ln u_i}{\partial \ln p_{i1}} = \frac{L_{i1}}{R_i} - \pi_{i1}.$$

If tourism increases the price in the tourist sector, locals will benefit if and only if the fraction of local labor employed in the tourist sector exceeds local's expenditure share on that sector. Keeping with the interpretation of the tourist sector as "restaurants", the intuition is straightforward: if increased tourism drives up the prices in restaurants, that's good news for waitstaff, but bad news for patrons.

3.2 A general theoretical framework for local analysis

The simple model above conveys the simple take-away that whether or not tourism is good for locals depends on how tourism (a) changes local prices; and (b) whether locals are net producers or consumers of the goods for which prices changed. We now extend the framework above to incorporate a rich geography of production and consumption across a city.

Suppose that a city (Barcelona) comprises a set of N city blocks, which we index by n (for the location of residence) and i (for the location of production and/or consumption). City blocks are separated by spatial frictions; we assume these spatial frictions are ad valorem and exogenous but otherwise the geography is completely general. The presence of spatial frictions are meant to capture the fact (verified below) that all else equal, it is more costly for residents to commute or purchase goods from city blocks the further from their residence. Suppose there are $S + 1$ sectors, where as above we assume sector $s = 0$ is a costlessly traded sector whose numeraire price is fixed by the broader economy outside of Barcelona. We view each location \times sector pair as a distinct good. Inhabitants of the city make three choices: (1) where to live; (2) which good to produce; and (3) how much of each good to consume.

For both production and consumption, we retain similar assumptions to the simple framework above. Production of good (i, s) combines labor L_{is} and a specific factor M_{is} in a constant returns to scale production function $Q_{is} = F_{is}(L_{is}, M_{is})$. Residents' consumption

choices are made to maximize a homothetic indirect utility function $u_n = \frac{v_n}{G(\{p_{nis}\}_{s \in \{0,1\}})} B_n$ with some price aggregator $G(\cdot)$, where v_n is the income of residents of block n and B_n is an (exogenous) amenity value of residing in location n .

A general expression for a change in welfare

We now proceed to derive a general expression for the impact of any shock on the welfare of residents. As above, we can fully (log) differentiate the indirect utility function and apply Shepherd's lemma to yield the change in utility from any change in prices and incomes:

$$d \ln u_n = \partial \ln v_n - \sum_{i,s} \pi_{nis} \partial \ln p_{is}, \quad (3)$$

where π_{nis} is the expenditure share of residents of block n on sector s in produced in location i . (Note that the price changes are not n specific: this is a result of the assumption that spatial frictions are ad valorem). The intuition for equation (3) is the same as for equation (1): the welfare impact of any shock depends on (1) its impact on resident's income; and (2) an expenditure share weighted impact on the prices of all goods.

Like above, we can also derive a relationship between the change in income of individuals residing in location n and the change in wages across all sectors and locations in the city. Suppose that residents of location n allocate their labor across the production of goods in order to maximize their income, i.e. they solve:

$$v_n = \max_{\{L_{nis}\}} \sum_{i,s} w_{nis} L_{nis} \text{ s.t. } H_n(\{L_{nis}\}) = R_n,$$

where $w_{nis} \equiv w_{is}/\mu_{nis}$ for some spatial friction $\mu_{nis} \geq 1$ and $H(\cdot)$ is a (weakly) convex function that captures potential decreasing returns to scale to allocating more labor to the production of a particular product, which we refer to as a "commuting function", and R_n is the labor endowment of locals residing in location n (which we hold fixed in the analysis throughout). Then applying the envelope theorem to this choice yields:

$$d \ln v_n = \sum_{i,s} \sigma_{nis} \partial \ln w_{is},$$

where $\sigma_{nis} \equiv \frac{w_{nis} L_{nis}}{v_n}$ is the share of income residents in n earn from good (i, s) . If we further assume that wages are equalized across sectors within location and that the spatial frictions incurred in commuting from n to i are the same for all sectors (which together imply

$w_{nis} = w_{ni}$), this expression simplifies to:

$$d \ln v_n = \sum_i \sigma_{ni} \partial \ln w_i, \quad (4)$$

where $\sigma_{ni} \equiv \frac{w_{ni} L_{ni}}{v_n}$ is the share of income residents in n earn from producing in location i . Equation (4) is intuitive: it says that a change in wages in any location i affects the income of residents in location n proportional to the share of income residents in n earn in i . Substituting in equation (4) into equation (3) yields the following expression for the change in utility of residents of location n to any change in wages and prices throughout the city:

$$d \ln u_n = \sum_i \sigma_{ni} \partial \ln w_i - \sum_{i,s} \pi_{nis} \partial \ln p_{is}. \quad (5)$$

Equation (5) forms the basis of the analysis that follows. It shows that there are two necessary ingredients in order to determine the welfare impact of any shock (including tourism) on local residents. The first necessary ingredient is knowledge of existing income shares (i.e. $\{\sigma_{ni}\}$) and expenditure shares (i.e. $\{\pi_{nis}\}$). Fortunately for us, we have access to detailed commuting flow data that can be used to reconstruct income shares and detailed expenditure data that records the location of the residence as well as the location and sector of the purchase to reconstruct expenditure shares.

The second necessary ingredient is knowledge of how a shock changes wages (across locations within the city) and prices (across all goods, i.e. locations \times sectors, within the city). While in principal these wage and price changes are estimable, in practice it is infeasible to simultaneously estimate N distinct wage changes and $N \times S$ distinct price changes. In what follows, we pursue two complementary strategies to overcome this limitation: first, in results presented in Section 4.1, we empirically estimate average price and wage elasticities across all locations in the city, which allows us to recover an average welfare impact of tourism; second, we impose simple equilibrium market clearing conditions to derive expressions for all wage and price changes throughout the city as a function of observed data. We turn to these derivations next.

A general expression for the (short-run) price and wage effects of a tourism shock

The results thus far have relied solely on the optimization on the part of residents; as a result, we have not needed to impose any general equilibrium conditions for the economy as a whole. Here, we show that imposing standard market clearing conditions allow us to trace out the impact of a tourism shock on prices and wages throughout the city. We model

tourism as expenditure E^T by non-residents on goods produced in Barcelona; we now derive how an (exogenous) increase in E^T – a tourism shock – affects prices and wages throughout Barcelona. For the time being, we hold labor allocations and expenditure shares fixed, which we refer to as the “short-run.” This is perhaps an appropriate assumption given that our empirical context examines impacts of tourism by comparing expenditure across months within a year; however, in the next section below we extend the framework to allow for adjustments to local labor allocations as well.

In all sectors $s \in \{1, \dots, S\}$ and for all locations $i \in \{1, \dots, N\}$, equilibrium prices ensure that supply equals demand. Let X_{nis} and X_{is}^T denote the expenditure on the good (i, s) by locals residing in n and tourists, respectively. Then this market clearing condition requires:

$$p_{is}Q_{is} = \sum_n X_{isn} + X_{is}^T$$

Re-writing expenditure as a function of incomes and expenditure shares yields and substituting in the resident’s labor allocation choice yields:

$$p_{is}Q_{is} = \sum_n \pi_{nis} \sum_l L_{nl} \mu_{nl} w_l + \pi_{is}^T E^T, \quad (6)$$

which is our first equilibrium condition.

Our second equilibrium condition relates the labor income to the total income derived for each good. Define $\beta_{is} \equiv \frac{w_{is} L_{is}}{p_{is} Q_{is}}$ as the labor income share, so that $p_{is} Q_{is} = \frac{1}{\beta_{is}} w_{is} L_{is}$. Maintaining the assumption that wages are equalized across sectors within location allows us to sum equation (6) across sectors to write total labor returns at the location level as a function of total demand in the location:

$$w_i L_i \left(\sum_s \frac{1}{\beta_{is}} \left(\frac{L_{is}}{L_i} \right) \right) = \sum_s \left(\sum_n \pi_{nis} \sum_l L_{nl} \mu_{nl} w_l + \pi_{is}^T E^T \right). \quad (7)$$

Equations (6) and (7) together define $N \times (S \times 1)$ system of equations that can be written as a function of the $N \times S$ endogenous prices $\{p_{is}\}_{s=\{1, \dots, S\}}^{i \in \{1, \dots, N\}}$ and N endogenous wages $\{w_i\}_{i \in \{1, \dots, N\}}$. Implicitly differentiating the system – holding constant labor allocations, labor income shares, and expenditure shares – yields:

$$\frac{\partial \ln p_{is}}{\partial \ln E^T} = \frac{X_{is}^T}{y_{is}} + \sum_n \frac{v_n}{y_{is}} \pi_{nis} \sum_j \sigma_{nj} \frac{\partial \ln w_j}{\partial \ln E^T} \quad (8)$$

$$\left[\frac{\partial \ln w_i}{\partial \ln E^T} \right]_i = (\mathbf{I} - \mathbf{A})^{-1} \left[\frac{\sum_s X_{is}^T}{\sum_s y_{is}} \right]_i, \quad (9)$$

where $\mathbf{A} \equiv \left[\sum_s \sum_n \pi_{nis} \frac{\nu_n}{y_{is}} \sigma_{nj} \right]_{ij}$. Equations (8) and (9) trace the short-run impact of an increase on tourist spending on prices and wages throughout the city. The expressions are intuitive. Consider first the price elasticity expression. The direct impact of a tourist shock is to increase the demand for good (i, s) proportional to the initial tourist expenditure share in that good; as quantities produced do not respond in the short-run, this results in a corresponding increase in prices. However, a tourist shock also has an indirect impact on demand (and hence prices) by changing local resident's income (and hence their demand) as well. The extent to which demand from local resident's changes depends on how both how much their income changes (captured by $\sum_j \sigma_{nj} \frac{\partial \ln w_j}{\partial \ln E^T}$ term) and how much they spend on a particular good (the $\sum_n \frac{\nu_n}{y_{is}} \pi_{nis}$ term).

Consider now the wage elasticity expression. Re-expressing the Leontief inverse in a Neumann series yields a more intuitive expression:

$$\frac{\partial \ln w_i}{\partial \ln E^T} = \frac{\sum_s X_{is}^T}{\sum_s y_{is}} + \sum_j \sum_s \sum_n \pi_{nis} \frac{\nu_n}{y_{is}} \sigma_{nj} \left(\frac{\sum_s X_{js}^T}{\sum_s y_{js}} \right) + \dots \quad (10)$$

Equation (10) shows that the direct impact of a tourism shock on wages is to increase demand for goods produced in location, which it does in proportion to tourist's initial share of expenditure. However, tourist shocks also have indirect effects on demand (and hence wages). The first degree indirect effect is that some of the the direct impact on wages elsewhere translate in changes in demand; for example, a local resident who earns additional tourist income (the direct effect) will spend that dollar elsewhere in the city (the first degree effect). The first degree indirect effect then in turn generates a second degree indirect effect (as another resident paid by the first resident spends that additional income elsewhere), and the process repeats ad infinitum, converging to the expression in equation (9).

There are two key take-aways from equations (8) and (9). First, how exactly a tourism shock changes prices and wages throughout the city depend on the interaction of the spatial patterns of income and consumption of local residents. Second, given knowledge of these spatial patterns (along with knowledge of the spatial pattern of consumption by tourists), one can apply equations (8) and (9) to determine the price and wage impacts of a tourism shock. This, in turn, can be combined with equation (5) to allows us to determine the (short-run, first-order) welfare impacts of tourism solely as function of observed data, an approach we pursue below in Section 4.2.1.

3.3 A specific theoretical framework for global analysis

The advantage of the framework presented in the previous subsection is its generality: the expressions derive hold for any homothetic preferences, any constant returns to scale production functions, and any commuting function. The disadvantage of the framework is that the welfare expressions are valid only for small tourist shocks and they hold only when expenditure shares and labor allocations are held constant (the “short run”). In this section, we present a complementary framework where we pursue the opposite tactic: we assume particular set of preferences and production functions and then derive welfare expressions that hold for arbitrarily sized tourist shocks and account for changes in labor allocations.

Setup

The setup is a special case of the previous section. On the production side, we assume Cobb-Douglas production functions for each sector:

$$Q_{is} = A_{is} L_{is}^{\beta_s} K_{is}^{1-\beta_s}.$$

On the consumption side, we assume consumers have nested CES preferences, with the outer nest across sectors (with elasticity η) and the inner nest across locations within sector with elasticity σ_s , which generates the following indirect utility function for residents in n :

$$u_n = \frac{v_n}{\left(\sum_{s=0}^S \alpha_s \left(\left(\sum_{i=1}^N \gamma_{is} \tau_{nis}^{1-\sigma_s} p_{is}^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}} \right)^{1-\eta} \right)^{\frac{1}{1-\eta}}} B_n, \quad (11)$$

where $\tau_{isn} \geq 1$ is an iceberg trade cost, α_s is an exogenous sector specific preference shifter, and γ_{is} is an exogenous good-sector preference shifter. This in turn generates the following demand function for residents:

$$X_{nis} = \left(\frac{\gamma_{is} \tau_{isn}^{1-\sigma_s} p_{is}^{1-\sigma_s}}{\sum_j \gamma_{js} \tau_{jns}^{1-\sigma_s} p_{js}^{1-\sigma_s}} \right) \left(\frac{\alpha_s \left(\left(\sum_{i=1}^N \gamma_{is} \tau_{nis}^{1-\sigma_s} p_{is}^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}} \right)^{1-\eta}}{\sum_{s=0}^S \alpha_s \left(\left(\sum_{i=1}^N \gamma_{is} \tau_{nis}^{1-\sigma_s} p_{is}^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}} \right)^{1-\eta}} \right) v_n. \quad (12)$$

and for tourists:

$$X_{is}^T = \left(\frac{\gamma_{is}^T p_{is}^{1-\sigma_s}}{\sum_j \gamma_{js}^T p_{js}^{1-\sigma_s}} \right) \left(\frac{\alpha_s^T \left(\left(\sum_{i=1}^N \gamma_{is}^T p_{is}^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}} \right)^{1-\eta}}{\sum_{s=0}^S \alpha_s^T \left(\left(\sum_{i=1}^N \gamma_{is}^T p_{is}^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}} \right)^{1-\eta}} \right) E^T, \quad (13)$$

where γ_{is}^T and α_s^T are the tourist goods and sector preference shifters. Finally, on the commuting side, we continue to assume that workers are perfectly mobile within location (so that wages are equalized across sectors in equilibrium) and furthermore assume the commuting function H_n is such that:

$$v_n = \left(\sum_i \mu_{ni}^{-\theta} w_i^\theta \right)^{\frac{1}{\theta}}, \quad (14)$$

which in turn generates the following expression for the number of workers residing in n and working in i :

$$L_{ni} = \frac{\mu_{ni}^{-\theta} w_i^\theta}{\sum_i \mu_{ni}^{-\theta} w_i^\theta} R_n.$$

An expression for the change in welfare

We now derive an expression for the change in local welfare in response to any shock to the economy. Relative to the derivation of equation (5) in Section 3.2, this shock need not be small, but the derivation does rely on the particular functional form assumptions. To do so, we employ the “exact hat” notation of [Dekle, Eaton, and Kortum \(2008\)](#) and [Costinot and Rodriguez-Clare \(2013\)](#) where $\hat{x} = x'/x$ indicates the ratio of a variable in its post-shock state (x') relative to its initial pre-shock state (x). Substituting the commuting income equation (14) into the indirect utility function, taking ratios between pre- and post-shock variables, and re-arranging using expressions for consumption and commuting shares yields:

$$\ln \hat{u}_n = \frac{1}{\theta} \ln \left(\sum_i \sigma_{ni} \hat{\mu}_{ni}^{-\theta} \hat{w}_i^\theta \right) + \ln \hat{B}_n - \frac{1}{1-\eta} \ln \left(\sum_{s=0}^S \pi_{n,s} \hat{\alpha}_s \left(\left(\sum_{i=1}^N \pi_{nis} \hat{\gamma}_{is} \hat{r}_{nis}^{1-\sigma_s} \hat{p}_{is}^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}} \right)^{1-\eta} \right),$$

where $\sigma_{ni} \equiv \frac{L_{ni}}{L_n}$ is the share of residents residing in n who work in location i , $\pi_{n,s} \equiv \frac{\sum_i X_{isn}}{\sum_{i,s} X_{isn}}$ is the sectoral expenditure share of residents residing in n on all goods from sector s , and $\pi_{isn} \equiv \frac{X_{isn}}{\sum_i X_{isn}}$ is the within-sector expenditure share of residents in n on goods from location i in sector s . Focusing on shocks that like an increase in tourism expenditure that do not

affect commuting costs, trade costs, preferences, or amenities, this expression simplifies to:

$$\ln \hat{u}_n = \frac{1}{\theta} \ln \left(\sum_i \sigma_{ni} \hat{w}_i^\theta \right) - \frac{1}{1-\eta} \ln \left(\sum_{s=0}^S \pi_{n,s} \left(\left(\sum_{i=1}^N \pi_{nis} \hat{p}_{is}^{1-\sigma_s} \right)^{\frac{1}{1-\sigma_s}} \right)^{1-\eta} \right). \quad (15)$$

Equation (15) extends the expression of the local welfare changes presented in equation (5) to show how welfare changes in response to arbitrarily large changes in wages and prices under the assumption of constant elasticity preferences and commuting patterns. The basic intuition of how price and wage changes affect utility remains unchanged: just as in equation (5), welfare effects can be separated into an income effect (which depends on observed commuting patterns) and an aggregate of price effects (which depend on observed expenditure shares). For large shocks, however, the expression also depends on how elastic labor supply and consumer demand is to price changes (i.e. the θ , η , and $\{\sigma_s\}$ parameters).

An expression for the price and wage effects of a tourism shock

We now derive an expression for how wages and prices change in response to a tourist shock. As in Section 3.2, we do so by imposing standard general equilibrium market clearing conditions. Unlike Section 3.2, however, we allow both labor allocations and expenditure shares to endogenously respond to the tourist shock.

As in Section 3.2, we begin with the following market clearing condition, which holds for all $s \in \{1, \dots, S\}$ (i.e. all sectors except the costlessly tradable sector $s = 0$) and all locations $i \in \{1, \dots, N\}$:

$$p_{is} Q_{is} = \sum_n X_{nis} + X_{is}^T.$$

Expressing the system in changes and using the above expressions for the income of tourists, the expenditure of locals and tourists, and the fact that workers' wage is equal to their marginal product times the output price (which allows us to express output as a function of

wages and prices, i.e. $Q_{is} = A_{is}\beta_s^{\frac{\beta_s}{1-\beta_s}} \left(\frac{p_{is}}{w_i}\right)^{\frac{\beta_s}{1-\beta_s}} K_{is}$, we have:

$$\begin{aligned} \hat{p}_{is}^{\frac{1}{1-\beta_s}} \hat{w}_i^{-\frac{\beta_s}{1-\beta_s}} &= \sum_n \left(\frac{X_{nis}}{y_{is}}\right) \frac{\left(\left(\sum_{i=1}^N \pi_{nis} \hat{p}_{is}^{1-\sigma_s}\right)^{\frac{1}{1-\sigma_s}}\right)^{1-\eta}}{\sum_{s=0}^S \left(\left(\pi_{n,s}\right) \left(\left(\sum_{i=1}^N \pi_{nis} \hat{p}_{is}^{1-\sigma_s}\right)^{\frac{1}{1-\sigma_s}}\right)^{1-\eta}\right)} \frac{\hat{p}_{is}^{1-\sigma_s}}{\sum_j \pi_{jns} \hat{p}_{js}^{1-\sigma_s}} \left(\sum_i \sigma_{ni} \hat{w}_i^\theta\right)^{\frac{1}{\theta}} \\ &+ \frac{X_{is}^T}{y_{is}} \frac{\left(\left(\sum_{i=1}^N \pi_{is}^T \hat{p}_{is}^{1-\sigma_s}\right)^{\frac{1}{1-\sigma_s}}\right)^{1-\eta}}{\sum_{s=0}^S \left(\pi_s^T \left(\left(\sum_{i=1}^N \pi_{is}^T \hat{p}_{is}^{1-\sigma_s}\right)^{\frac{1}{1-\sigma_s}}\right)^{1-\eta}\right)} \frac{\hat{p}_{is}^{1-\sigma_s}}{\left(\pi_{js}^T\right) \hat{p}_{js}^{1-\sigma_s}} \hat{E}^T, \end{aligned} \quad (16)$$

where $\pi_{is}^T \equiv \frac{X_{is}^T}{\sum_i X_{is}^T}$ and $\pi_s^T \equiv \frac{\sum_i X_{is}^T}{\sum_{i,s} X_{is}^T}$ are tourist's within and across sector expenditure shares. Equation (16) extends the derivation from equation (8) to large shocks and no longer holds labor allocations and expenditure shares fixed. While the two expressions are similar, note that the left hand side of (16) depends on the labor share in the production function (as this determines the extent to which wages change as labor reallocates its production across sectors within location) and that the right hand side incorporates both the demand elasticities η and $\{\sigma_s\}$ (as expenditure shares may adjust to large shocks) and the labor supply elasticity θ (as local residents may reallocate their labor across locations).

The second market clearing condition is a labor market clearing condition, which requires that the labor employed across sectors within location is equal to the total amount of labor flowing from all residences to that location:

$$\sum_{s=0}^S L_{is} = \sum_{n=1}^N L_{ni}.$$

(Note that in Section 3.2, this labor market clearing condition was automatically satisfied as we held labor allocations fixed). Multiplying both sides by wages, using again the fact that workers' wage is equal to their marginal product times the output price, and expressing in changes yields:

$$\sum_s \left(\frac{\beta_s y_{is}}{\sum_s \beta_s y_{is}}\right) \hat{p}_{is}^{\frac{1}{1-\beta_s}} \hat{w}_i^{-\frac{\beta_s}{1-\beta_s}} = \sum_n \sigma_{ni} \left(\frac{R_n w_i}{\sum_s \beta_s y_{is}}\right) \frac{\hat{w}_i^{1+\theta}}{\sum_j \sigma_{nj} \hat{w}_j^\theta}. \quad (17)$$

Equations (16) and (17) can be solved jointly to determine all $S \times N$ price changes and N wages changes in response to any change in tourist expenditure \hat{E}^T , accounting for the endogenous reallocation of labor and expenditure (which we refer to as the ‘‘long run’’ ef-

fect). As in the local derivation presented in Section 3.2, the observed spatial distribution of expenditures (for both locals and tourists) and bilateral commuting flows play a crucial role in determining these wage and price changes; now, however, knowledge of supply elasticities $\{\beta_s\}$ and θ and demand elasticities $\{\sigma_s\}$ and η are also necessary. Given these price and wage changes, equation (15) then determines the welfare impact of this change in tourist expenditure for all residents in the city. We use such a strategy to determine the long-run welfare consequences of tourism for locals in Section 4.2.2.

4 Is tourism good for locals? Three empirical approaches

Armed with the data described in Section 2.1 and the theoretical frameworks presented in Section 3, we now turn to an empirical analysis of the welfare impacts of tourism on locals in Barcelona. Like in the previous section, we pursue three complementary empirical approaches. In the first approach, we seek to estimate average impacts of tourism on prices and wages across the city using a deductive fully empirical approach. We next build an inductive model-based framework grounded in the theoretical underpinnings described in the previous section. Finally, we combine both worlds in a hybrid deductive-inductive approach that allows us to get heterogeneous effects and, at the same time, test whether the theory is able to capture the heterogeneity observed in the data.

4.1 A deductive (regression based) approach

We first pursue an empirical regression-based approach to identify the average elasticities of wages and sector-specific prices. We then combine these estimated average elasticities in a very general welfare expression to calculate the welfare impact of tourism. This analysis has the advantage, under the appropriate identification assumptions, of providing estimates that are independent from the theoretical structure. The downside is that we will be able to identify only average price effects and wage effects across all locations (and hence average welfare effects). We begin by discussing the exogenous identification sources and introducing our Bartik-type instrument.

A Bartik instrument exploiting seasonal variation in tourist origin composition and sector spending patterns

Our goal is to estimate the local price and wage effects of tourism. A generic regression of price or wages on tourist activity, however, would be inappropriate, as there may be correlated preference shocks between tourists and locals. We therefore build an instrument

to identify β . We use our structural counterpart of tourist demand to develop a Bartik instrument for tourist spending. Intuitively, we rely on two facts (see Stylized Fact 1): (1) tourists from different countries of origin allocate their expenditure differently within Barcelona; and (2) the composition of tourists from different countries of origin in Barcelona changes throughout the year.

Inspired by our data, we consider different kinds of tourists, denoted by g (domestic and foreign),² each of whom has a distinct demand structure for local products, given by equation (13). Total tourist spending in a given location can thus be expressed as:

$$X_i^T = \sum_g E_g^T \times \Phi_{ig}, \quad (18)$$

where E_g^T is the total city-wide spending by group g , $\Phi_{ig} \equiv \sum_s \frac{\gamma_{isg}^T p_{is}^{1-\sigma_c}}{\sum_j \gamma_{jgs}^T p_{js}^{1-\sigma_s}} \kappa_{sg}^T$, and κ_{sg}^T is used for exposition purposes to summarize the CES-implied expenditure share in equation (13). Taking the sources of exogenous variation to be the total spending of each group, group-specific consumption shares, and preferences, we totally differentiate equation (18) to obtain an explicit expression for the sources of exogenous variation in terms of initial shares and group-specific spending variation:

$$\begin{aligned} \frac{dX_i^T}{X_i^T} &= \sum_g \frac{E_g^T \Phi_{ig}}{X_i^T} \frac{dE_g^T}{E_g^T} + \sum_g \sum_s \frac{X_{isg}^T}{X_i^T} \frac{d\kappa_{sg}^T}{\kappa_{sg}^T} + \sum_g \sum_c \frac{X_{isg}^T}{X_i^T} \frac{d\gamma_{isg}^T}{\gamma_{isg}^T} \\ &= \sum_g \underbrace{\frac{X_{ig}^T}{X_i^T}}_{\varsigma_{ig|i}^T} \frac{dE_g^T}{E_g^T} + \sum_g \sum_s \underbrace{\frac{X_{isg}^T}{X_i^T}}_{\varsigma_{isg|i} = \varsigma_{isg|ig} \varsigma_{ig|i}} \frac{d\kappa_{sg}^T}{\kappa_{sg}^T} + \sum_g \sum_c \frac{X_{isg}^T}{X_i^T} \frac{d\gamma_{isg}^T}{\gamma_{isg}^T} \end{aligned}$$

Taking it to the data, we construct changes in location i 's tourist expenditures $g_i^T \equiv \frac{\Delta E_i^T}{E_i^T}$ as:

$$g_i^T = \underbrace{\sum_g \varsigma_{i,g|i} \times g_{E_g}^T}_{\text{Group Composition}} + \underbrace{\sum_g \sum_s \varsigma_{i,s,g|i} \times g_{\kappa,s,g}^T}_{\text{Seasonal Demand}}, \quad (19)$$

where $g_{E_g}^T \equiv \frac{\Delta E_g^T}{E_g^T}$ and $g_{\kappa,sg}^T = \frac{\Delta \kappa_{sg}^T}{\kappa_{sg}^T}$ denote changes in total group's income and in within-group category spending, respectively; the $\varsigma_{i,g|i} \equiv \frac{E_{ig}^T}{E_i^T}$ captures the initial (or average) tourist

²In ongoing work, we refine this categorization to distinguish between foreign tourists from different countries of origin.

group composition of tourist spending in i , and $\varsigma_{i,s,g|i} \equiv \frac{E_{i,s,g}^T}{E_i^T}$ combines the initial group composition $\varsigma_{i,g|i}$ and the initial within-group consumption shares.³

We use (19) to define different instruments that exploit group composition $\sum_g \varsigma_{i,g|i} \times g_{Eg}^T$, seasonal demand changes $\sum_g \sum_s \varsigma_{i,s,g|i} \times g_{\kappa,s,g}^T$, or both $\sum_g \varsigma_{i,g|i} \times g_{Eg}^T + \sum_g \sum_s \varsigma_{i,s,g|i} \times g_{\kappa,s,g}^T$. Our current benchmark results are calculated using the group composition IV alone. We define the *initial* shares to be orthogonal to seasonal demand shocks by building seasonal-averages of the shares over the full period available in our data.

Price Regressions

We first examine the impact of tourism on local prices. We measure changes in local prices using observed changes in resident's expenditure shares. We first regress expenditure shares of locals residing in location n on goods (i, s) in month m in year t on a set of origin-product-month-year and destination-product-month-year fixed effects, as well as the distance between them:

$$\log \pi_{nismt} = \phi_s \log \tau_{ni} + \log \gamma_{nstm} + \log \delta_{istm} + u_{nismt}.$$

To identify the scale of the fixed effects, we choose as the numeraire the car industry, which we assume accords to the tradable sector in the theory. The car industry is an appropriate empirical analog to the tradable sector as it is likely that automobile prices are set on the world market; moreover, both SEAT and Nissan manufacturer cars in Barcelona, so local employment in the automobile sector is substantial. Consistent with equation (12), we then measure changes in local prices as follows:

$$\Delta \ln \mathbf{p}_{ismt} \equiv \frac{1}{1 - \sigma_s} \Delta \ln \delta_{istm},$$

where Δ indicates a first difference (taken across months) and the calibration of σ_s is discussed below. (Note the particular choice of σ_s simply scales the estimated coefficients that follow).

We then regress changes in local prices of goods in sector s in location i on changes in tourist expenditure in the same location, instrumenting tourist expenditure with the Bartik

³Notice that the unobserved confounder is given by the weighted aggregate of the preference component:

$$\epsilon_i = \sum_g \sum_s \pi_{isg} \times g_{isg}^\gamma.$$

instrument introduced in Subsection 4.1:

$$\Delta \ln p_{ismt} = \gamma_{is} + \gamma_{ts} + \beta_s^p \times \Delta \log E_{itm}^T + \epsilon_{ismt}, \quad (20)$$

where the location-sector fixed effects γ_{is} and the year-sector fixed effects γ_{ts} capture unobserved time trends at the location-sector and year-sector level, respectively. From equation (12), we know that the residual ϵ_{ismt} includes any time varying changes in local preferences for a good (i, s) , so the exclusion restriction requires that aggregate variation in the composition of tourists is uncorrelated with location-specific changes in local preferences within Barcelona, conditional on fixed effects.

Figure 10 presents the results of regression (20) for all sectors. We find that tourism increased prices in all sectors, a result that is statistically significant at conventional levels in 18 (of 19) sectors. The effects are substantial, with a 10% increase in tourist expenditure increasing prices by 1% - 3%, depending on the sector.

Wage Regressions

We now examine the impact of tourism on local wages. Wages are not directly observed in our data, but can be recovered using observed local expenditure data and commuting patterns as follows:

$$w_{imt} = \sum_{n=1}^N \left(\frac{L_{ni}}{R_n} \right) v_{nmt},$$

where local expenditures serve as an empirical proxy for v_{nmt} and data on commuting patterns inform the variables L_{ni} and R_n . There are a few caveats to the wage imputation procedure. First, we assume away all savings and behavioral responses (such as income smoothing) to income fluctuations on discretionary credit-card spending. Locals immediately spend their income on discretionary spending. In ongoing work, we are assessing the veracity of this assumption using an alternative measure of income based on direct deposits of paychecks. Second, we assume our commuting data is representative and reflective of the population covered by the spending data. As mentioned above, we verify the accuracy of the commuting data by cross-referencing across several different commuting datasets.

Proceeding analogously to the previous section, we can estimate the wage impact of tourism by regressing changes in wages in location i in month m in year t on tourist expenditure:

$$\Delta \ln w_{imt} = \gamma_{it} + \gamma_{im} + \gamma_{tm} + \beta^w \times \Delta \log E_{itm}^T + \epsilon_{imt}, \quad (21)$$

instrumenting with the Bartik instrument introduced in Subsection 4.1.

Column 1 of Table 1 presents the results. We find that a 10% increase in tourist expenditure increases wages by 0.5%, with a 95% confidence interval of 0.2% to 0.9%.

Is tourism good for the locals (on average)?

Combining the price and wage estimates, we now assess whether or not tourism is good for locals on average using welfare equation (5). Since the price and wage effects are averaged across all locations, the spatial patterns of commuting do not affect the welfare estimates and the spatial patterns of consumption affect the welfare estimates only inasmuch as locals differ in their sectoral expenditure shares. Evaluating the price effects at the aggregate (city-wide) sectoral expenditure shares, we find that the price index elasticity to tourism expenditure is 0.18 to 0.28. Since the wage elasticity to tourism expenditure is 0.02 to 0.09, this means that the average welfare impact of tourism was -0.16 to -0.19. To put these number in context, consider the increase in tourist expenditures between February and July, on average. Using our data, this number is 70.3%. Using the average for our elasticities above—-0.23 and 0.05, respectively—, this seasonal increase would translate into an income welfare gain of 3.5% and a price-index welfare deterioration of 16.17%, with a net welfare deterioration of 12.67% for residents on average. This average effect turns out to mask substantial heterogeneity in welfare effects across residents of Barcelona, which we now turn to examining.

4.2 An inductive (theory based) approach

In our deductive approach, we rely on reduced form elasticities to estimate the effect of tourists on locals, without relying on a fully specified general-equilibrium model. We now complement the above analysis with a inductive, theory based, approach. We do so in two, complementary ways. First, in subsection 4.1, we consider the short run, using equations (8) and (9) and only relying on basic market clearing conditions. This simulates the short run effect of tourism on prices and incomes, holding constant labor allocations and expenditure shares. Second, in subsection 4.2, we adopt structure for both supply and demand from (3.3) and consider the “long run” impacts of tourist shocks, allowing for both the reallocation of labor and consumption.

To utilize our inductive approach, we bolster the above credit card data with three additional data sources. First we re-weight resident expenditures across categories using the Spanish consumer expenditure survey. Second, we add housing rental rates by neighborhood (Barrio) from Idealista (as housing search tool). Lastly, we balance the share of tourist spending in Barcelona to 20%, reflecting national accounting statistics. This provides us

with a comprehensive approximation of consumer spending in Barcelona.

4.2.1 Is tourism good for locals (in the short run)?

We first solve equations (8) and (9) using expenditure, commuting, and income data to determine the price and welfare effects of tourist shocks in the short-run.

We present the results in Figure (11). Consistent with the deductive results, the inductive results predict that tourism causes both prices and incomes to rise, with the effects concentrated near the city center where tourism expenditure is greatest. The magnitudes of the effects are also quite similar, with both price index elasticities varying from 0.1 to 0.3 and the wage elasticities varying from 0.05 to 0.5. The net effect for most residents is very close to zero; with the positive income effect relatively stronger in the city center and the price effect relatively stronger in the city periphery.

Recall from Section 3.2 that these short run elasticities hold for general preferences and production structures and therefore do not require the estimation of any structural parameters. Moreover, unlike the deductive method, we can estimate heterogeneous welfare effects across the city. However, there are two disadvantages. First, this is a local estimate of price, income and utility effects and can only simulate small shocks in the short run. For example workers cannot change commuting patterns, nor can consumption patterns change. Second, this method heavily relies on market clearing conditions, with labor and product markets efficiently clearing. We now address the first limitation by allowing for larger longer-run changes. We then consider a hybrid deductive-inductive approach that addresses the second concern.

4.2.2 Is tourism good for locals (in the long run)?

We build on the “exact hat-algebra” approach of [Dekle, Eaton, and Kortum \(2008\)](#) to consider longer run changes to the spatial equilibrium of Barcelona following a tourism shock. In particular, Equations (16) and (17) can be solved jointly to determine price wages changes in response to any change in tourist expenditure \hat{E}^T . This method allows for both the reallocation of labor, as well as expenditures for both locals and tourists (which we refer to as the “long run” effect).

To determine welfare effects in the previous estimate, looking at the short run, we simply needed data and market clearing conditions. This estimate is both more computationally demanding, as well as requiring estimates of the factor share of labor $\{\beta_s\}$ and θ and demand elasticities $\{\sigma_s\}$ and η . Given these price and wage changes, equation (15) then determines the welfare impact of this change in tourist expenditure for all residents in the city.

This approach relies on us taking a stance on both the production and demand functions. In terms of our fundamentals governing supply and demand, we currently calibrate them using the spatial literature and leveraging macro-economic data.⁴ For the factor share of labor, β_s , we use macro-economic labor shares for Spain across 20 economic sectors, averaging 66%. For labor supply θ , we use estimates from the [Monte, Redding, and Rossi-Hansberg \(2018\)](#), with $\theta = 3.3$. For demand parameters σ_s , we simply use the average from [Hottman, Redding, and Weinstein \(2016\)](#) of 3.9. For the upper-level substitution elasticity η , we follow [Hobijn and Nechio \(2019\)](#) and use a median estimate of 1.8.

We present results for a 50% increase in tourism (from the 2018 baseline) in Figure (12). In panel (a), we show the income effect, with the elasticity of the residential price index to the tourist shock. In panel (b), we show the price effect, with the elasticity of wages to a tourist shock. In panel (c), we combine the results and show the aggregate welfare effect. In particular, there is wide heterogeneity. While there is a small positive average effect, there are regions where higher prices outweigh even a large, positive wage effect. In particular, locations with local consumption, but commuters into the central city, disproportionately benefit. At the other extreme, residents in the central city, see a much larger price effect, while only experiencing a similar wage effect. This is a reversal of the short term effects, highlighting the role played by the commuting and consumption elasticities.

4.3 A hybrid deductive-inductive approach

The advantage of the deductive approach of Section 4.1 was that it did not require any theoretical assumptions to identify the average price and wage effects of tourism across the city. The advantage of the inductive approach of Section 4.2 was that allowed us to estimate the heterogeneous price and wage effects throughout the city (and hence the heterogeneous welfare impacts on tourists). In this section, we propose a hybrid deductive-inductive approach that seeks to combine the best of both approaches by allowing the data to determine the heterogeneous wage and price effects throughout the city.

The methodology

The approach is simple and straightforward: we adapt the basic specification of Section 4.1 to allow for heterogeneous treatment effects, where the source of the heterogeneity is given by the distribution of the model-implied elasticities calculated in Section 4.2. Consider first the price regressions. Let $\eta_{is}^p \equiv \frac{\partial \log p_{is}}{\partial \log X_i^T}$ be the model-predicted elasticity of the price in location i and sector s to an increase in tourism expenditure. We use the sector-specific distribution of

⁴In ongoing work, we estimate the model parameters using the tourist shock as identifying variation.

η_{is}^p over locations to categorize location i as high or low elasticity. In particular, we use the median as a threshold above which a location is considered high elasticity, within a sector s ; intuitively a high-elasticity location indicates the theory predicts that location i will have a larger than median price response in sector s . We then interact the low and high- η_{is}^p indicators—denoted by $\mathbb{1}_{is}^{p,low} \equiv \mathbb{1}\{\eta_{is}^p \leq \text{median}(\eta_{is}^p | s)\}$ and $\mathbb{1}_{is}^{p,high} \equiv \mathbb{1}\{\eta_{is}^p > \text{median}(\eta_{is}^p | s)\}$, respectively—with the tourist expenditure in price regression (20), i.e:

$$\Delta \ln p_{ismt} = \gamma_{is} + \gamma_{ts} + \beta_s^{p,low} \times \mathbb{1}_{is}^{p,low} \times \Delta \log E_{imt}^T + \beta_s^{p,high} \times \mathbb{1}_{is}^{p,high} \times \Delta \log E_{imt}^T + \epsilon_{ismt}, \quad (22)$$

where $\beta_s^{p,HTE} \equiv \beta_s^{p,high} - \beta_s^{p,low} > 0$ indicates that locations where the model predicts the price response should be larger indeed have larger price responses, i.e. the theory is correct.

Similarly, let $\eta_i^w \equiv \frac{\partial \log w_i}{\partial \log X_i^T}$ be the model-predicted elasticity of the wage in location i to an increase in tourism expenditure. We can then interact the low and high- η_i^w indicators— $\mathbb{1}_i^{w,low} \equiv \mathbb{1}\{\eta_i^w \leq \text{median}(\eta_i^w)\}$ and $\mathbb{1}_i^{w,high} \equiv \mathbb{1}\{\eta_i^w > \text{median}(\eta_i^w)\}$ —with tourist expenditure in wage regression (21), i.e:

$$\Delta \ln w_{imt} = \gamma_i + \gamma_t + \beta^{w,low} \times \mathbb{1}_i^{w,low} \times \Delta \log E_{imt}^T + \beta^{w,high} \times \mathbb{1}_i^{w,high} \times \Delta \log E_{imt}^T + \epsilon_{imt}, \quad (23)$$

where $\beta^{w,HTE} \equiv \beta^{w,high} - \beta^{w,low} > 0$ also indicates that locations where the model predicts the wage response should be larger indeed have larger wage responses, i.e. the theory is correct.

Armed with estimates of $\beta_s^{p,high}, \beta_s^{p,low}, \beta^{w,high}, \beta^{w,low}$, we can construct heterogeneous empirical price elasticities $\frac{\partial \ln p_{is}}{\partial \ln E_i^T} = \beta_s^{p,low} \times \mathbb{1}_{is}^{w,low} + \beta_s^{p,high} \times \mathbb{1}_{is}^{w,high}$ and empirical wage elasticities $\frac{\partial \ln w_i}{\partial \ln E_i^T} = \beta^{w,low} \times \mathbb{1}_{is}^{w,low} + \beta^{w,high} \times \mathbb{1}_{is}^{w,high}$ that can then be combined with observed spatial consumption and income patterns to construct the welfare impact of tourism for estimates of residents in each locations within the city using equation (5). In what follows, we employ this approach for three different model predicted elasticities.

The “zero-degree” elasticities

We begin with the simplest measure of heterogeneity in price and wage elasticities. The first term in both the short run elasticity derivations for prices in equation (8) and for wages in equation (10) is the fraction of tourism expenditure on the good and in the location, respectively. Intuitively, the greater the share of tourism expenditure, the larger the impact an increase in tourism on demand for a good (or a location) and the greater the price (and wage) effect. This first term captures only the “zero-degree” effect of tourism – it abstracts from how this change in demand affects demand (and prices) for goods elsewhere as the

workers producing in that location use their increase in income on goods elsewhere. As a short-run elasticity, it also abstracts from any change in labor allocations or expenditure. However, it is intuitive and readily observable, so we begin here. Define the the “zero-degree” model predicted price elasticity $\eta_{is}^{p,0} \equiv \frac{E_{is}^T}{E_{is}^R + E_{is}^T}$ as the tourist expenditure share of a good (i, s) and the “zero-degree” model predicted wage elasticity as $\eta_i^{w,0} \equiv \frac{\sum_s E_{is}^T}{\sum_s (E_{is}^R + E_{is}^T)}$ as the of a location i .

Figure 13 depicts the the regression results of the price regression (22). For this case, the identified heterogeneous effects differ from those given by the zero-degree elasticities for most sectors. Only in 7 out of the (non-normalized) 19 sectors the effect on high $\eta_{is}^{p,0}$ is larger than on low $\eta_{is}^{p,0}$. In other words, locations where the model predicts the price response should be larger do not necessarily have larger price responses. Notice, however, that the pooled regression that exploits sectoral variation as well does predict a larger effect on locations with high model-implied elasticities. That is, the aggregate (across sectors) price response reflects heterogeneous effects that are in line with those in the zero-degree measure of elasticity. Column 2 of Table 1 presents the results of the wage regression. In contrast to the price regression, we do not estimate substantial heterogeneity, with the interaction coefficient with the tourist share of a location very close to zero.

The short-run elasticities

We now present the results for the short run elasticity derivations for prices in equation (8) and for wages in equation (10), incorporating all the higher degree terms. Figure 14 depicts the the regression results of the price regression by comparing the estimated elasticity for locations with high and low $\eta_{is}^{p,shortrun}$ for each sector. Incorporating the higher degree terms has a significant impact on the estimated heterogeneous effects and the relationship with the model-implied elasticities. In those case, 12 out of 19 sectors feature a larger price response in those locations with higher short-run price elasticities. In other words, in a majority of sectors, locations where the model predicts the price response should be larger do have larger price responses. Column 3 of Table 1 presents the results of the wage regression. Looking at the short-run effect, there is a very small positive effect in locations with higher short-run wage elasticities.

The long-run elasticities

Finally, we present the results for the long run elasticity derivations for prices calculated by solving equations (16) and (17) simultaneously. Figure 15 depicts the the regression results of the price regression by comparing the estimated elasticity for locations with high and low

$\eta_{is}^{p,longrun}$ for each sector. As is evident from looking at this figure, the results are very similar, though somehow dampened, to those in the short-run case. In other words, in a majority of sectors, locations where the model predicts the price response should be larger do have larger price responses. Column 4 of Table 1 presents the results of the wage regression. Looking at the long-run effect, there is a substantial positive effect in locations with higher long-run wage elasticities. Translating our results to percentiles, moving from the 15th to the 86th percentile of locations, a 10% increase in tourist spending increases wages from 0.35% to 1.1%.

Implied welfare effects

Figure 16, 17, 18 depict the implied welfare effects from equation (5) using the price and wage elasticities constructed using the three methods discussed in this subsection 4.3, as well as the observed spatial consumption and income patterns. Inspired by the hybrid estimation results discussed above, we take the short-run elasticities and Figure 17 as a benchmark. We find that, on average, local workers suffer slightly from tourism, but these average effects mask substantial heterogeneity across space. In particular, the median of the welfare changes distribution over locations is a *negative* 12.3 percent, but the effects range from a negative 19.5 percent to a positive 3.6 percent when looking at the 10th and 90th percentiles, respectively. Looking at the spatial patterns, the welfare effects are a result of price and income patterns, where the price effects feature an clear inner-outer city pattern and the income effects, in addition, reflect the cross-neighborhoods income inequality. In particular, we see that residents in the city center and those near tourist locations bear the largest price changes but also enjoy substantial income gains. In contrast, residents of peripheral neighborhoods suffer lower but still sizable price changes, with the income gains varying between different outer city locations: some (northern periphery) experience none and some (southern periphery) get moderate income benefits from tourism.

5 Conclusion

In this paper we have developed quantitative theoretical and empirical tools to analyze the welfare impacts of a very particular urban demand shocks - tourism. At the heart of the analysis is the combination of two different elements: Firstly, a novel expenditure database that describes in fine geographical detail the spatial consumption patterns of locals as a function of their residential location across time and sectors and secondly, a novel quantitative workhorse model - an urban Ricardo-Viner model - that allows us to trace out the rich price effects across location and sectors in response to an external demand shock, which we use

as our foundation to develop complementary empirical tools. While we apply the data and methodology to a particular question, we believe that the framework can be generally applied to phenomena that shift the effective demand in an urban context and might generate rich price effects with distributional impacts (e.g. urban renewal or gentrification). The urban Ricardo-Viner model in conjunction with expenditure data can be flexibly adjusted to any empirical context and offers a rich theoretical and empirical playground. In that sense this framework might be a first step towards closing the gap between efforts in economics and recent innovations in urban planning to create 'digital twins' of cities, i.e. urban models that are constructed from rich spatio-temporal data and inform policy making - possibly in real-time.

Regarding the policy question at hand, we find that tourism on average hurts residents. More importantly, the average welfare numbers hide in plain sight a large degree of overall and spatial heterogeneity in the incidence. Because tourism is a spatially concentrated phenomenon, that - at least in the case of Barcelona - affects mostly the urban core, the price effects, mostly driven by a severe price index deterioration for residents with a high expenditure share across the urban core, tend to be concentrated amongst a smaller share of the population. This insight should reshape the discussion of tourism, and should highlight to what extent urban policy must consider questions regarding spatial equity.

References

- AGARWAL, S., J. B. JENSEN, AND F. MONTE (2017): “Consumer Mobility and the Local Structure of Consumption Industries,” NBER Working Papers 23616, National Bureau of Economic Research, Inc.
- AHLFELDT, G. M., S. J. REDDING, D. M. STURM, AND N. WOLF (2015): “The economics of density: Evidence from the Berlin Wall,” *Econometrica*, 83(6), 2127–2189.
- ALLEN, T., C. ARKOLAKIS, AND Y. TAKAHASHI (2020): “Universal gravity,” *Journal of Political Economy*, 128(2), 393–433.
- ALMAGRO, M., AND T. DOMÍNGUEZ-IINO (2019): “Location Sorting and Endogenous Amenities: Evidence from Amsterdam,” .
- ATHEY, S., D. BLEI, R. DONNELLY, F. RUIZ, AND T. SCHMIDT (2018): “Estimating Heterogeneous Consumer Preferences for Restaurants and Travel Time Using Mobile Location Data,” *AEA Papers and Proceedings*, 108, 64–67.
- CARVALHO, V., J. R. GARCIA, S. HANSEN, Á. ORTIZ, T. RODRIGO, S. RODRIGUEZ MORA, AND J. RUIZ (2020): “Tracking the COVID-19 crisis with high-resolution transaction data,” WorkingPaper 2020/16.
- COSTINOT, A., AND A. RODRIGUEZ-CLARE (2013): “Trade theory with numbers: Quantifying the consequences of globalization,” Discussion paper, National Bureau of Economic Research.
- COUTURE, V., J. DINGEL, A. GREEN, AND J. HANDBURY (2020): “Quantifying Social Interactions Using Smartphone Data,” .
- DAVIS, D. R., J. I. DINGEL, J. MONRAS, AND E. MORALES (2019): “How Segregated Is Urban Consumption?,” *Journal of Political Economy*, 127(4), 1684–1738.
- DEKLE, R., J. EATON, AND S. KORTUM (2008): “Global Rebalancing with Gravity: Measuring the Burden of Adjustment,” *IMF Staff Papers*, 55(3), 511–540.
- DIX-CARNEIRO, R., AND B. K. KOVAK (2017): “Trade Liberalization and Regional Dynamics,” *American Economic Review*, 107(10), 2908–46.
- DONALDSON, D. (2018): “Railroads of the Raj: Estimating the impact of transportation infrastructure,” *American Economic Review*, 108(4-5), 899–934.

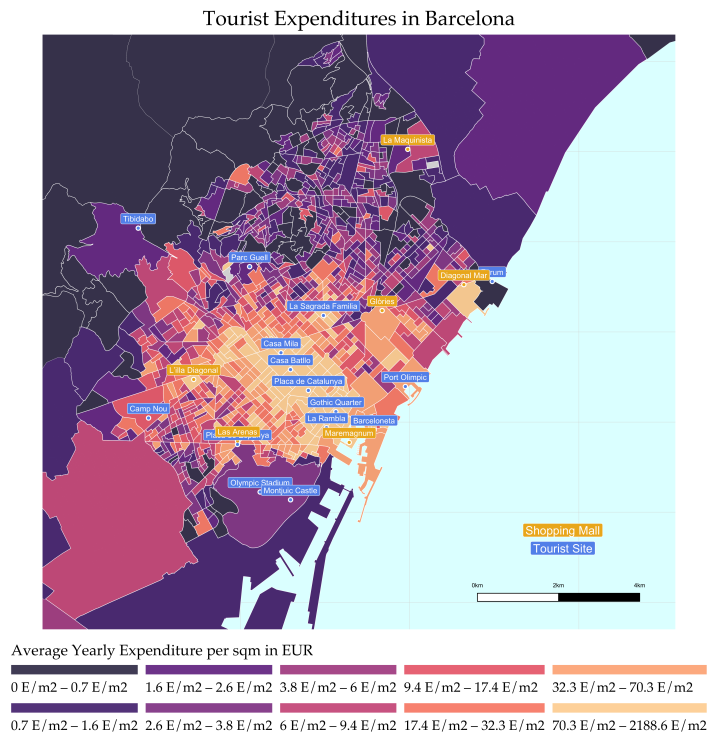
- FABER, B., AND C. GAUBERT (2019): “Tourism and Economic Development: Evidence from Mexico’s Coastline,” *American Economic Review*, 109(6), 2245–93.
- GARCÍA-LÓPEZ, M.-À., J. JOFRE-MONSENY, R. M. MAZZA, AND M. SEGÚ (2019): “Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona,” Working Papers 2019/05.
- HOBijn, B., AND F. NECHIO (2019): “Sticker shocks: using VAT changes to estimate upper-level elasticities of substitution,” *Journal of the European Economic Association*, 17(3), 799–833.
- HOTTMAN, C. J., S. J. REDDING, AND D. E. WEINSTEIN (2016): “Quantifying the sources of firm heterogeneity,” *The Quarterly Journal of Economics*, 131(3), 1291–1364.
- JONES, R. W. (1975): “Income distribution and effective protection in a multicommodity trade model,” *Journal of Economic Theory*, 11(1), 1–15.
- KOVAK, B. K. (2013): “Regional Effects of Trade Reform: What Is the Correct Measure of Liberalization?,” *American Economic Review*, 103(5), 1960–76.
- MONTE, F., S. J. REDDING, AND E. ROSSI-HANSBERG (2018): “Commuting, Migration, and Local Employment Elasticities,” *American Economic Review*, 108(12), 3855–90.
- MUSSA, M. (1974): “Tariffs and the Distribution of Income: The Importance of Factor Specificity, Substitutability, and Intensity in the Short and Long Run,” *Journal of Political Economy*, 82(6), 1191–1203.
- (1982): “Imperfect factor mobility and the distribution of income,” *Journal of International Economics*, 12(1-2), 125–141.

Figure 1: Signs in Barcelona



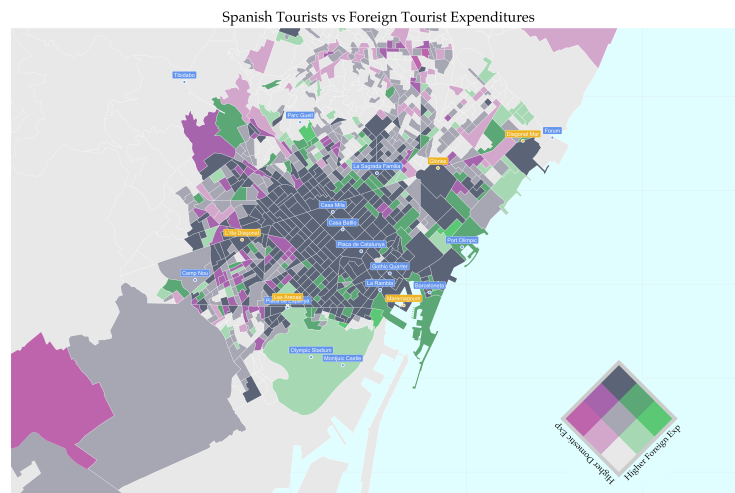
Notes: Source: “Why Barcelona Locals Really Hate Tourists.” *The Independent*. August 9, 2017.

Figure 2: Tourists spend disproportionately more in the city center




Notes: This figure shows the average yearly expenditure (normalized per square meter) in euros by tourists throughout the city of Barcelona.

Figure 3: Foreign tourists spend relatively more near the beach, whereas Spanish tourists spend relatively more near the malls



Notes: This map compares the spatial patterns of foreign and Spanish tourists, where expenditure is measured as total annual euros spent per square meter. Areas with high Spanish tourist expenditure and low foreign tourist expenditure are indicated in purple, whereas areas with high foreign tourist expenditure and low Spanish tourist expenditure are indicated in green. Expenditure measured in average monthly expenditure per square meter of the underlying tile. The boundary points for the tertiles are given by $([0,2.9],[2.9,14.4],[14.4,218.8])$ for tourist expenditure changes and by $([0,1.4],[1.4,4.9],[4.9,253.6])$ for local expenditure changes.

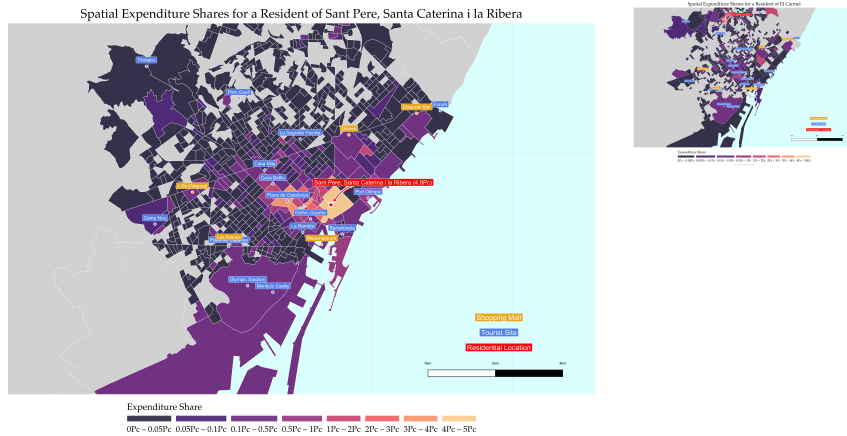
Figure 4: International tourist expenditure exhibits clear seasonality, Spanish tourists visit year round



Notes: This figure shows the total expenditure by month of locals (red), Spanish tourists (yellow), and international tourists (green). We also include expenditure by local non-Caixa bank customers (blue) and non-Caixa bank Spanish tourists (purple), where their residence is imputed as the location with their greatest expenditure throughout the year.

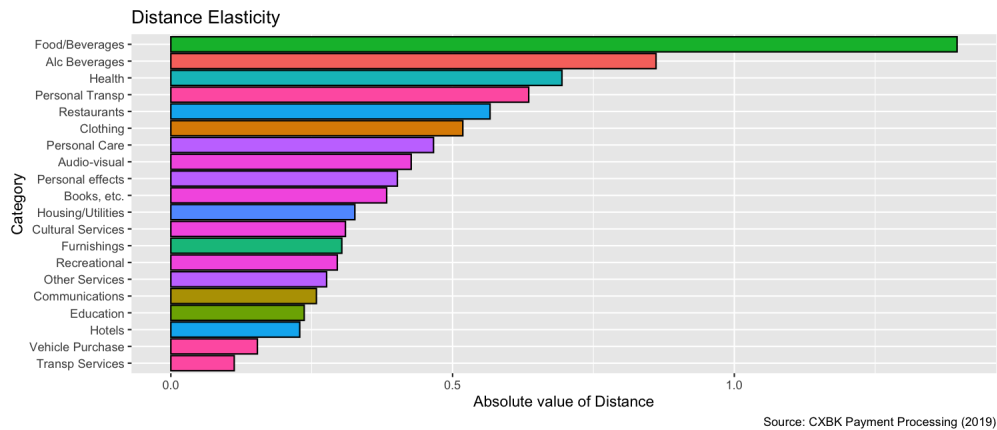
Figure 5: Locals spend more near their home

- (a) Expenditure shares for a local residing near the city center
- (b) Expenditure shares for a local residing far from the city center



Notes: This figure compares expenditure patterns for locals residing in different areas of the city. The left panel is the expenditure shares for a resident of Sant Pere, Santa Caterina i la Ribera (near the city center). The right panel is for El Carmel (far from the city center).


Figure 6: Residents spend more near their home, although the impact of distance is heterogeneous across sectors




Notes: This figure shows the impact of distance on expenditure by sector. The distance coefficient is estimated using sector-specific gravity regression of local expenditure shares on bilateral travel times with origin-sector-month and destination-sector-month fixed effects.

Figure 7: Residents are more likely to work nearby their home

(a) Commuting patterns for a resident near the city center

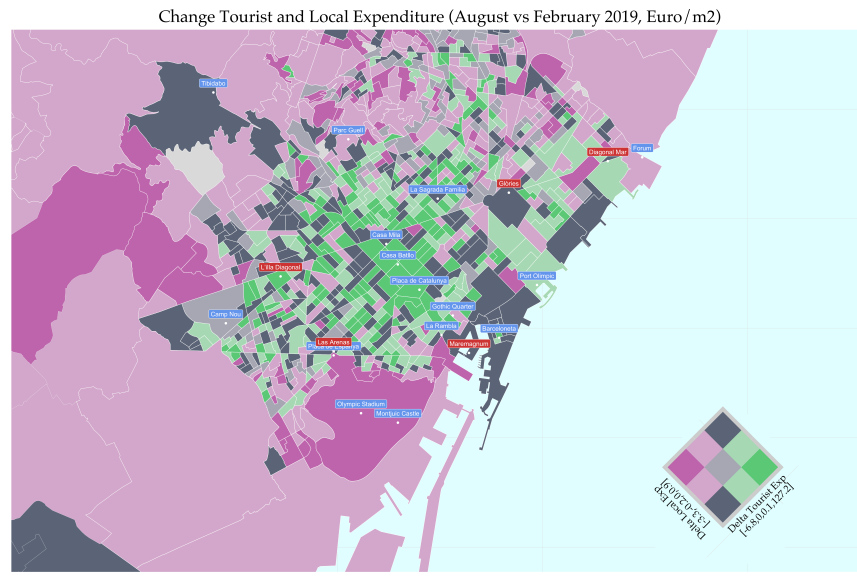


(b) Commuting patterns far from the city center



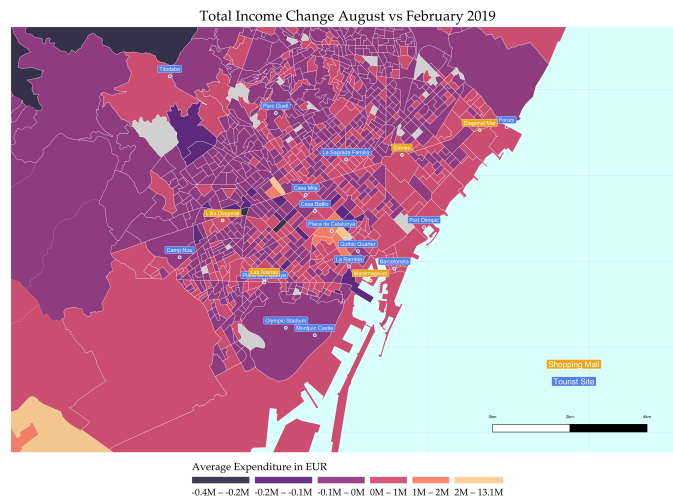
Notes: This figure compares the commuting patterns of locals residing in Poblenou near the city center (left panel) to those of residents residing far away from the city center in Carmel (right panel).

Figure 8: Tourist spending appears to crowd out and local expenditure in the high season



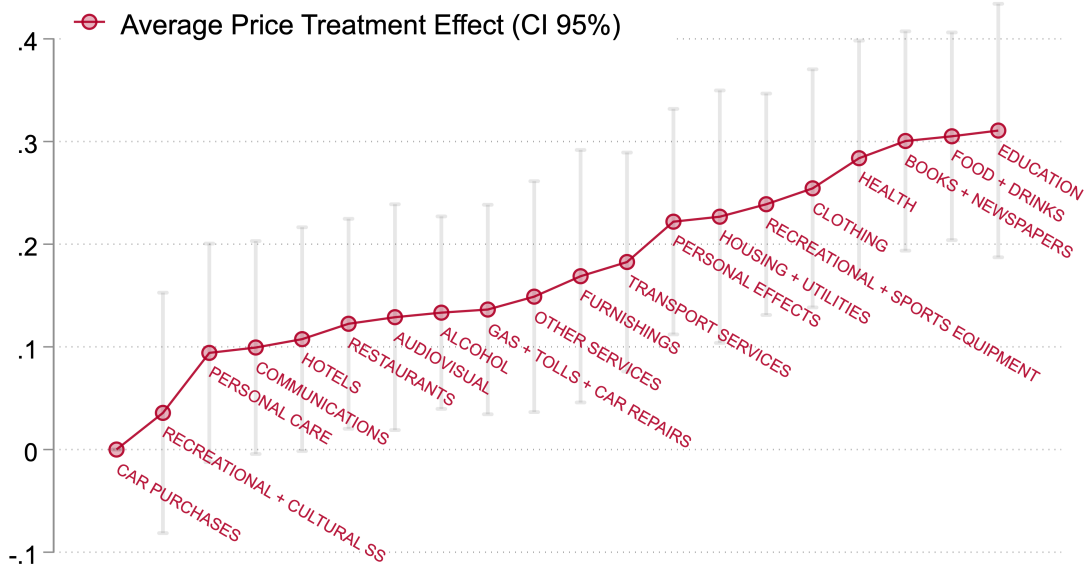
Notes: This figure compares the percentage change in tourist and local expenditure between the tourist high season (August) and the low season (February) in 2019. Locations where tourism expenditures increases and local expenditure decrease or grow weakly are demarked with green colors, while magenta colors mark the locations where the reverse is true. Expenditure changes are measured as level differences in average monthly expenditure per square meter of the underlying tile. Dark gray denotes locations where expenditures comove. The boundary points for the tertiles are given by $([-4.36, 0], [0, 0.09], [0.09, 19, 65])$ for tourist expenditure changes and by $([-0.9, -0.02], [-0.02, 0.01], [0.01, 1.13])$ for local expenditure changes.

Figure 9: Total sales increased more in touristy areas in the high season



Notes: This figure shows the change in total sales in the high season (August) versus the low season (February) in 2019.


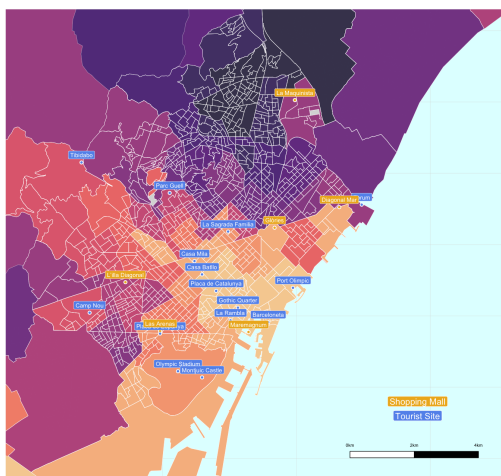
Figure 10: Tourism increased prices on average in all sectors




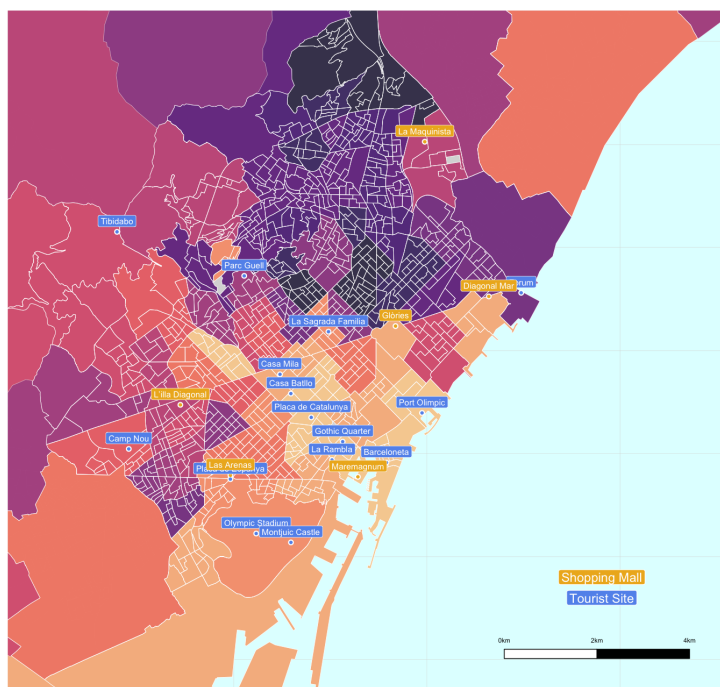
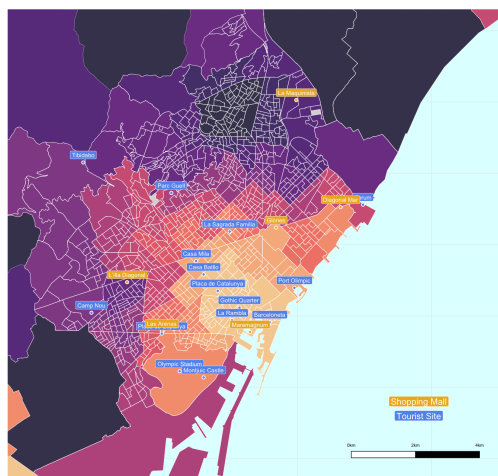
Notes: This figure shows the estimated price effects (and 95% confidence interval) of the price regression (20) for all sectors. The price in the car sector is normalized to one.

Figure 11: Inductive Approach: Is tourism good for locals in the short run?

(a) Income Effect



(b) Price Effect



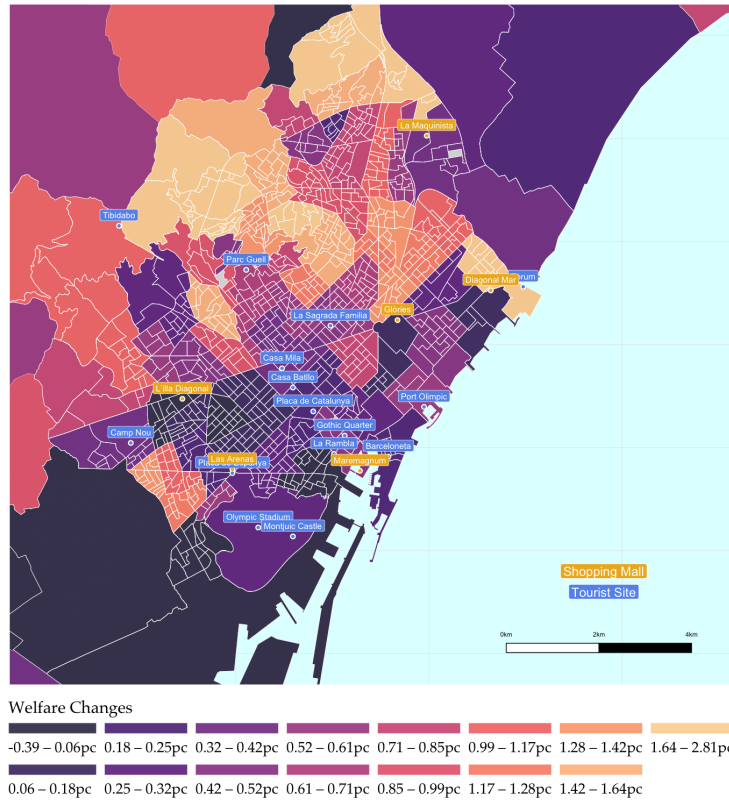
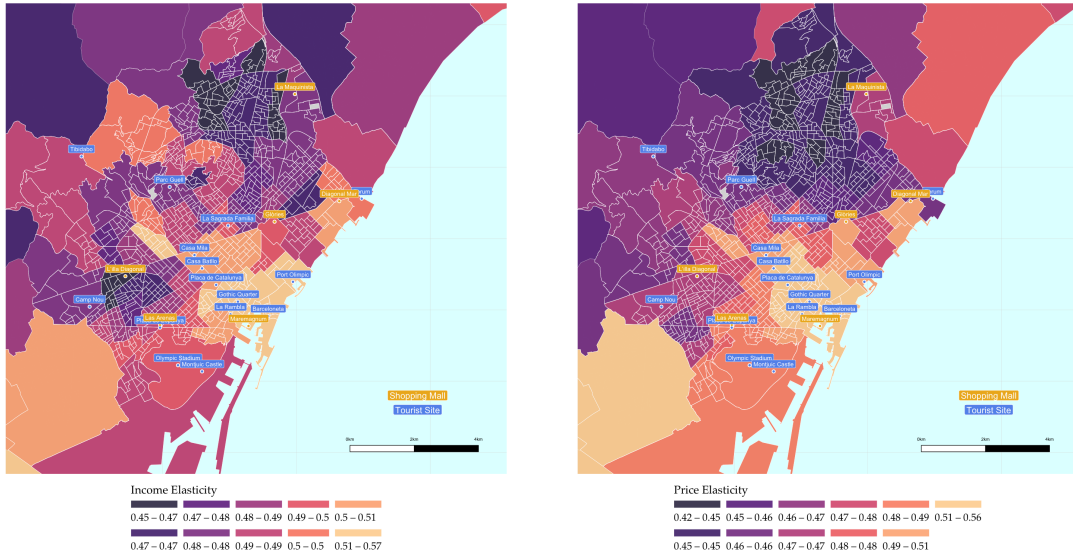
(c) Welfare Effect

Notes: This figure shows the “short-run” impact of tourism on wages (panel a), the price index (panel b), and total welfare (panel c), where the “short-run” holds constant local labor allocations and expenditure shares.

Figure 12: Inductive Approach: Is tourism good for locals in the long-run?

(a) Changes in Wages

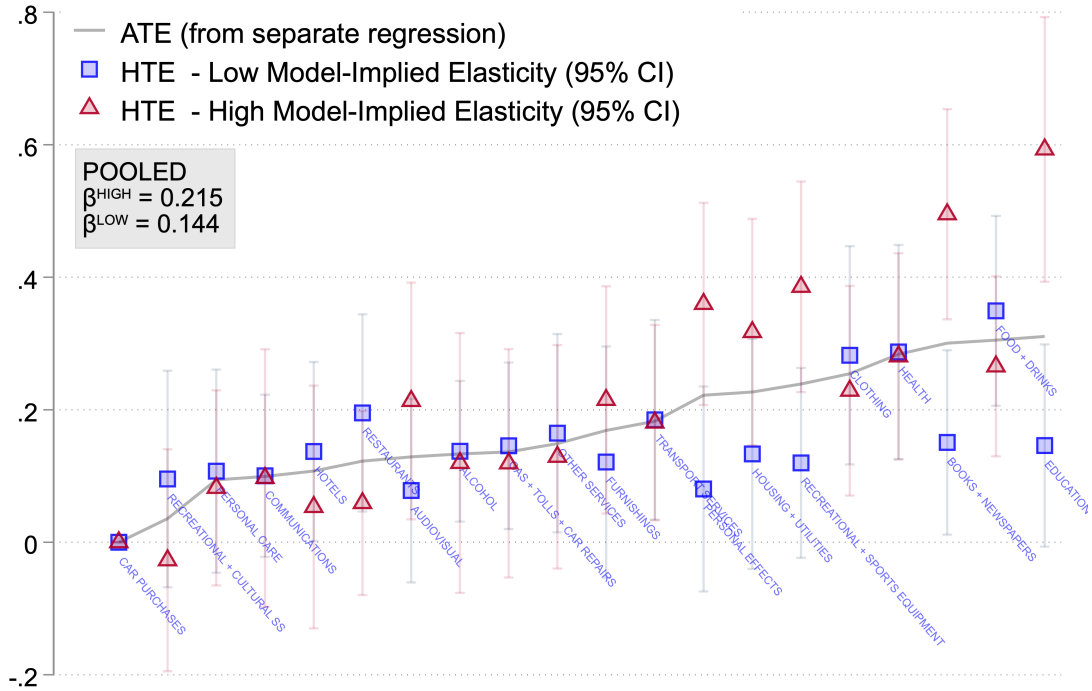
(b) Changes in Price Index



(c) Changes in Welfare


Notes: This figure shows the “long-run” impact of a 50% increase in tourism on wages (panel a), the price index (panel b), and total welfare (panel c), where by “long-run” we mean that local labor allocations and expenditure shares are able to adjust.

Figure 13: Hybrid Approach: Zero-degree elasticities



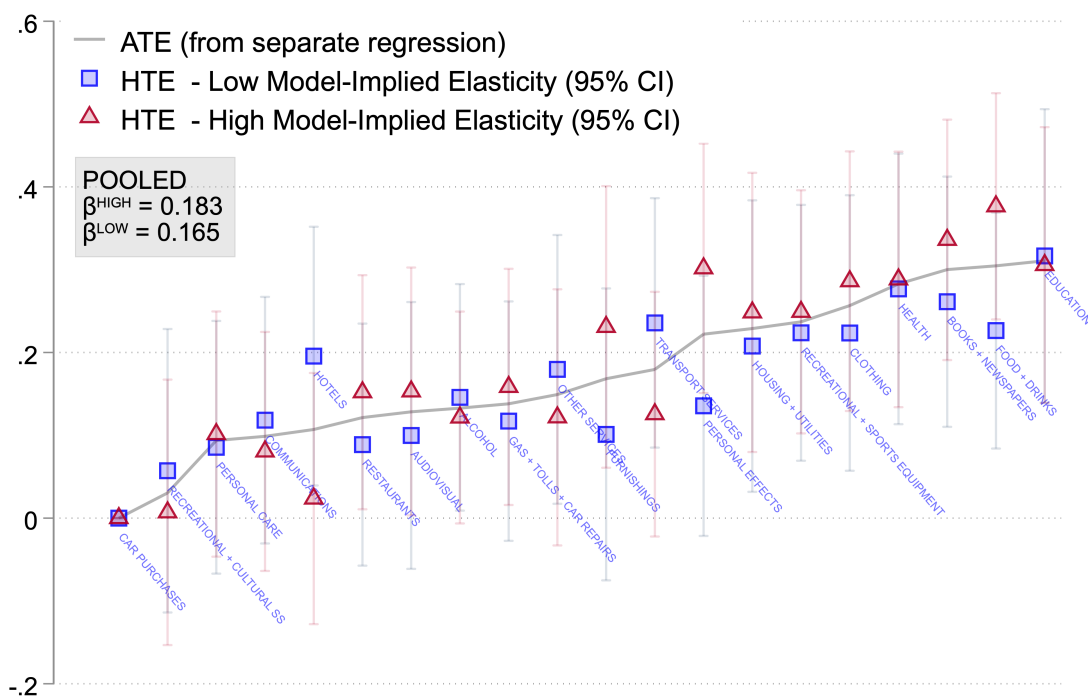
Notes: This figure shows the estimated total price effects (and 95% confidence interval) of the price regression (22) for all sectors. The solid, gray reference line corresponds to the ATE estimates from price regression (20) and is identical to the line in Figure 10. Low elasticity, in blue-diamond markers, identifies locations with model-implied elasticity lower or equal to the median. Accordingly, high elasticity, in red-triangle markers, identifies locations with model-implied elasticity higher than the median. The results of a regression of all sectors pooled together are included in a gray box below the legend. The price in the car sector is normalized to one.

Figure 14: Hybrid Approach: Short-run elasticities



Notes: This figure shows the estimated total price effects (and 95% confidence interval) of the price regression (22) for all sectors. The solid, gray reference line corresponds to the ATE estimates from price regression (20) and is identical to the line in Figure 10. Low elasticity, in blue-diamond markers, identifies locations with model-implied elasticity lower or equal to the median. Accordingly, high elasticity, in red-triangle markers, identifies locations with model-implied elasticity higher than the median. The results of a regression of all sectors pooled together are included in a gray box below the legend. The price in the car sector is normalized to one.


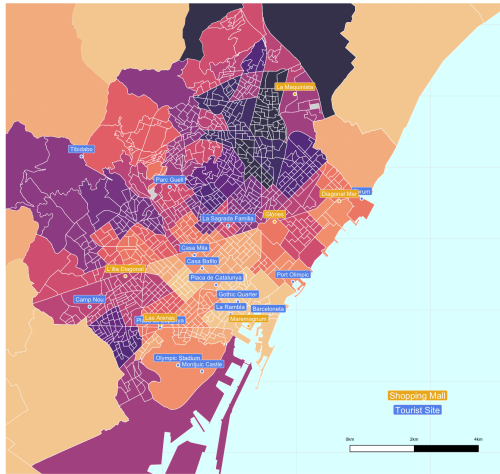
Figure 15: Hybrid Approach: Long-run elasticities



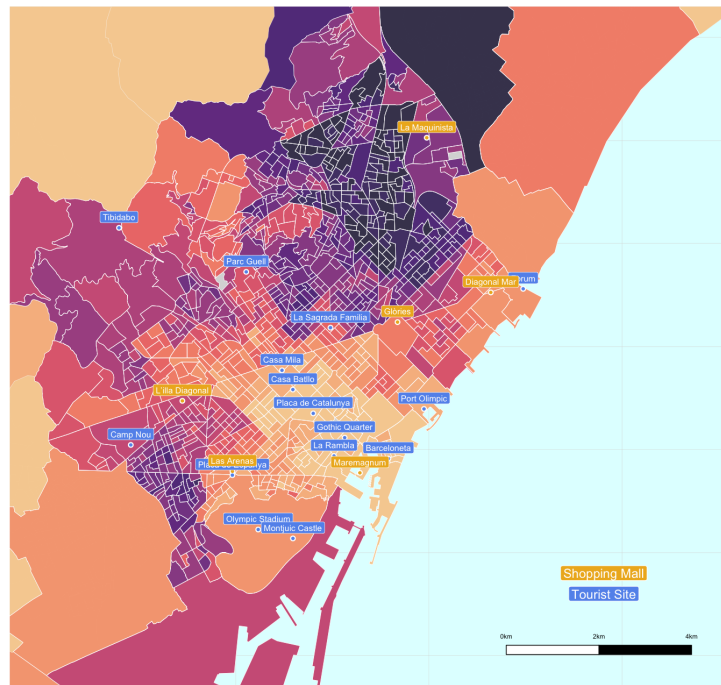

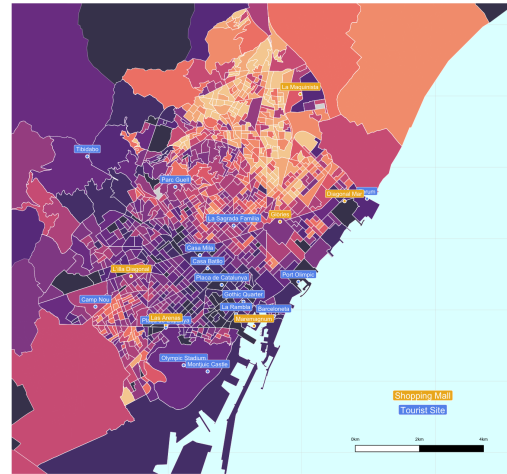
Notes: This figure shows the estimated total price effects (and 95% confidence interval) of the price regression (22) for all sectors. The solid, gray reference line corresponds to the ATE estimates from price regression (20) and is identical to the line in Figure 10. Low elasticity, in blue-diamond markers, identifies locations with model-implied elasticity lower or equal to the median. Accordingly, high elasticity, in red-triangle markers, identifies locations with model-implied elasticity higher than the median. The results of a regression of all sectors pooled together are included in a gray box below the legend. The price in the car sector is normalized to one.

Figure 16: Hybrid Approach: Is tourism good for locals - Degree-Zero?

(a) Income Effect



(b) Price Effect


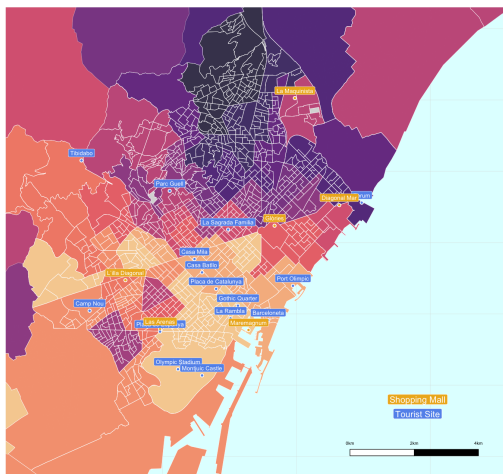


(c) Welfare Effect

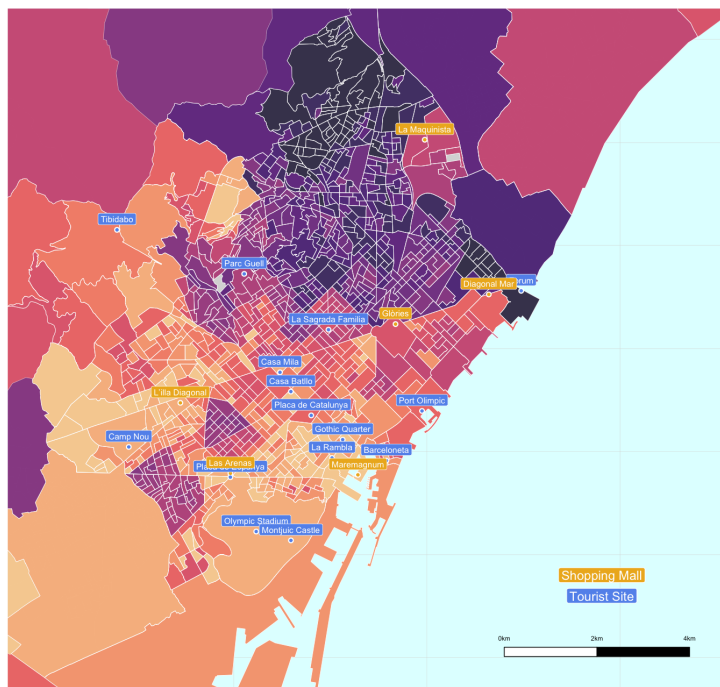
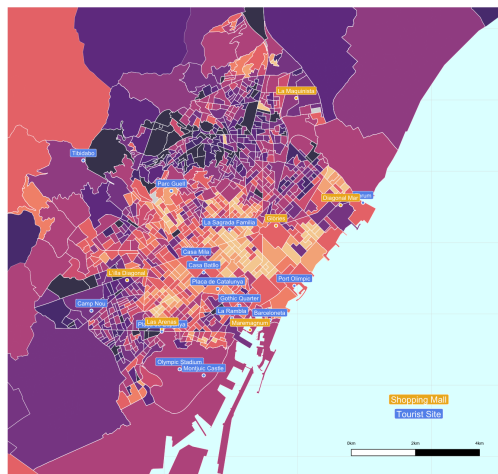
Notes: This figure shows the “short-run” impact of tourism on wages (panel a), the price index (panel b), and total welfare (panel c), where the “short-run” holds constant local labor allocations and expenditure shares.

Figure 17: Hybrid Approach: Is tourism good for locals in the short run?

(a) Income Effect



(b) Price Effect



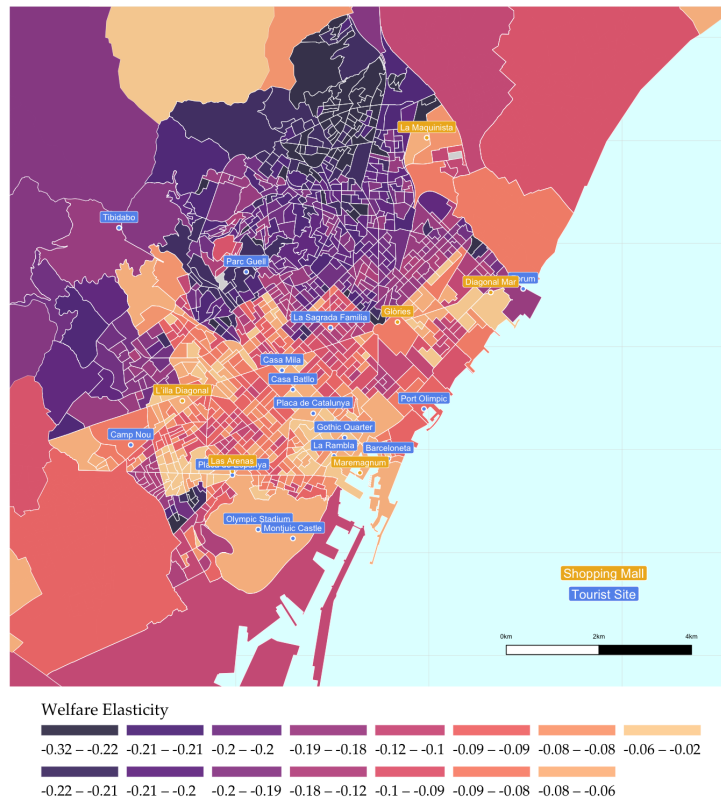
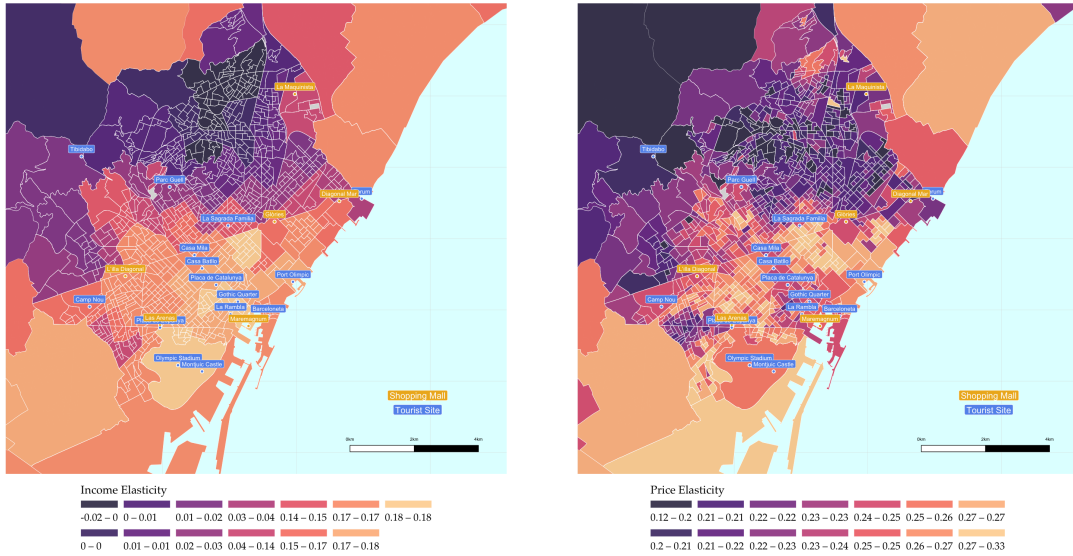
(c) Welfare Effect

Notes: This figure shows the “short-run” impact of tourism on wages (panel a), the price index (panel b), and total welfare (panel c), where the “short-run” holds constant local labor allocations and expenditure shares.

Figure 18: Hybrid Approach: Is tourism good for locals in the long-run?

(a) Changes in Wages

(b) Changes in Price Index



(c) Changes in Welfare

Notes: This figure shows the “long-run” impact of a 50% increase in tourism on wages (panel a), the price index (panel b), and total welfare (panel c), where by “long-run” we mean that local labor allocations and expenditure shares are able to adjust.

Table 1: Wage Regression Results

	(1)	(2)	(3)	(4)
	Baseline	Zero	SR	DEK
S.ln(Tourist Expenditures)	0.0530** (0.0173)	-0.0396 (0.0243)	0.00326 (0.0109)	-0.0232 (0.0165)
x Tourist Share > Median		0.193* (0.0822)		
x Short Run Wage Elasticity > Median			0.289** (0.0940)	
x Long Run Wage Elasticity > Median				0.212*** (0.0507)
Observations	24238	24238	24238	24238
IV	1	1	1	1
FE location-year	1	1	1	1
FE year-month-type	1	1	1	1
FE location-month	1	1	1	1

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Standard errors are clustered by location. The tourist share is the share of expenditures in a location accounted by tourist spending over the entire time period. The short run wage elasticities are computed using equation 9, and does not allow for the reallocation of consumption or workplaces. The long run wage elasticities are computed using equations 16 and 17, and account for the reallocation of consumption and workplaces. The level effects are absorbed by the location fixed-effects and omitted.

Table 2: Commuting Gravity Results

	(1)	(2)	(3)	(4)
	PPML	OLS	PPML	OLS
Log(Distance)	-4.628*** (0.313)	-2.121*** (0.138)		
Distance			-0.485*** (0.0294)	-0.127*** (0.0156)
Observations	11449	1633	11449	1633
FE: Origin	1	1	1	1
FE: Destination	1	1	1	1

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Standard errors are two-way clustered at origin and destination. Columns (1) and (3) are estimated using Pseudo-Poisson Maximum Likelihood, where commuting flows are: $E(\lambda_{ni}) = \exp(\alpha \log(d_{ni}) + \gamma_n + \delta_i)$. Columns (2) and (4) are estimated using Ordinary Least Squares, where commuting flows are: $\log(\lambda_{ij}) = \alpha \log(d_{ni}) + \gamma_n + \delta_i + \epsilon_{ni}$. Distances in minutes are computed using the simple average of transit times over commuting hours using a car and public transit. Travel times within a location is normalized to 2 minutes.

Table 3: Summary Statistics: Total Average Total Expenditure by month and time

	Local	Spanish Tourists	Foreign Tourists	Total
Total	1593.60 (54)	314.07 (11)	1062.589 (36)	2970.26
Jan	142.31 (63)	24.06 (11)	60.53 (27)	226.90 (100)
Feb	125.63 (59)	21.81 (10)	66.93 (31)	214.36 (100)
Mar	143.02 (58)	25.57 (10)	79.38 (32)	247.97 (100)
Apr	135.99 (52)	26.98 (10)	97.05 (37)	260.02 (100)
May	146.34 (53)	28.16 (10)	104.01 (37)	278.50 (100)
Jun	145.43 (53)	28.05 (10)	101.05 (37)	274.54 (100)
Jul	149.24 (50)	32.83 (11)	118.40 (39)	300.47 (100)
Aug	101.74 (41)	27.83 (11)	116.46 (47)	246.03 (100)
Sep	117.89 (49)	23.97 (10)	96.55 (40)	238.41 (100)
Oct	122.80 (51)	23.77 (10)	93.40 (39)	239.97 (100)
Nov	124.67 (57)	24.04 (11)	68.46 (32)	217.17 (100)
Dec	138.55 (61)	27.01 (12)	60.37 (27)	225.92 (100)

Notes: The table shows the average total expenditures (in million Euros) per month and across groups. The groups are aggregated to reflect our notion of locals (CXBK and non-CXBK customers), foreign tourists (transaction utilizing a credit card with a foreign issuer) and domestic Spanish tourists (cards that have their largest expenditure outside of the province of Barcelona).

Table 4: Summary Statistics: Total Average Total Expenditure 2-Digit COICOP

COICOP (2D)	COICOP (2D)	Local	Spanish Tourists	Foreign Tourists	Total	Survey (INE)	Survey Adj (INE)
11	Food/Beverages	32.82 (24.72)	1.32 (5.04)	4.51 (5.10)	38.66	12.96	23.82
21	Alc Beverages	1.97 (1.48)	0.07 (0.28)	0.60 (0.68)	2.64	0.71	1.31
31	Clothing	11.58 (8.72)	1.94 (7.39)	12.00 (13.55)	25.51	3.39	6.23
41	Housing/Utilities	2.81 (2.12)	0.78 (3.00)	0.59 (0.67)	4.19	5.33	9.80
51	Furnishings	10.03 (7.55)	3.32 (12.67)	2.01 (2.27)	15.35	0.88	1.62
61	Health	10.76 (8.10)	1.94 (7.40)	1.82 (2.06)	14.52	2.24	4.12
71	Vehicle Purchase	3.14 (2.36)	0.18 (0.67)	0.32 (0.36)	3.63	3.78	6.95
72	Personal Transp	7.27 (5.47)	2.06 (7.89)	0.70 (0.79)	10.03	6.38	11.73
73	Transp Services	10.13 (7.63)	6.52 (24.90)	9.61 (10.85)	26.26	1.90	3.49
81	Communications	0.30 (0.23)	0.02 (0.09)	0.08 (0.09)	0.40	0.33	0.61
91	Audio-visual	5.06 (3.81)	0.57 (2.17)	1.78 (2.01)	7.40	0.58	1.07
93	Recreational	2.62 (1.97)	0.27 (1.03)	1.21 (1.37)	4.09	1.43	2.63
94	Cultural Services	4.29 (3.23)	0.62 (2.38)	2.79 (3.15)	7.70	0.57	1.05
95	Books, etc	1.64 (1.23)	0.22 (0.85)	0.53 (0.60)	2.39	1.30	2.39
101	Education	1.11 (0.84)	0.10 (0.39)	0.61 (0.69)	1.82	0.77	1.41
111	Restaurants	17.73(13.35)	3.79 (14.46)	19.04 (21.50)	40.56	7.83	14.39
112	Hotels	1.13 (0.85)	1.49 (5.69)	23.12 (26.11)	25.75	1.21	2.22
121	Personal Care	4.84 (3.64)	0.32 (1.23)	0.97 (1.10)	6.14	2.53	4.65
123	Other	2.49 (1.88)	0.36 (1.37)	5.69 (6.42)	8.54	0.32	0.59
Total		131.72 (100)	25.88 (100)	87.97 (100)	245.58	54.4	100

Notes: The table shows the average total expenditures (in million Euros) per COICOP category and across groups. The groups are aggregated to reflect our notion of locals (CXBK and non-CXBK customers), foreign tourists (transaction utilizing a credit card with a foreign issuer) and domestic Spanish tourists (cards that have their largest expenditure outside of the province of Barcelona). We also report the corresponding expenditure share in the expenditure survey by INE for Catalonia. Since our consumption categories only add up to 54.4pc of total expenditures observed in the INE surveys, we construct an adjusted expenditure share measure from the surveys that accounts for this and is directly comparable to our expenditure shares.