Residential segregation and social segregation by race

Kirsten Cornelson

The most recent version of this paper can be found at:

http://individual.utoronto.ca/cornelson

Abstract

In this paper, I examine a determinant of social segregation by race in the United States: physical distance. Because U.S. cities are highly segregated, the time cost of interacting with a member of another race is typically higher than the cost of interacting with a same-race friend. The goal of this paper is to quantitatively assess the importance of this channel in explaining why people typically interact with members of their own race. Based on external estimates of consumers' costs of time, I simulate the frequency of cross-racial interactions that would occur if *only* distance mattered in determining individuals' choice of interaction partners. I compare the simulation results to a new measure of the actual frequency of inter-racial interactions based on Flickr photographs. I estimate that 25-30% of social segregation for whites in the U.S. is attributable to physical distance alone.

I would like to thank Dwayne Benjamin, Michael Baker and Aloysius Siow, as well as participants in several seminars at the University of Toronto, for their helpful comments on this paper.

1 Introduction

There is a high degree of physical and social isolation between blacks and whites in the United States. The statistics on residential segregation are well known. In 2010, the average black American lived in a Census block that was 54.1% black, despite the fact that blacks made up only 12.2% of the population as whole. Even more striking, Echenique and Fryer (2007) report that, as of 2000, over 60% of Census blocks in most states contained residents of only one race. Blacks and whites rarely marry each other: of all married couples with either a black or white spouse in 2014, just 1.2% were inter-racial. There is significant segregation in the workplace (Hellerstein and Neumark, 2008), in high school friendship networks within the same school (Echenique and Fryer, 2007) and between university classmates (Marmaros and Sacerdote, 2006).

To what extent are the physical and social dimensions of racial isolation causally related to each other? Sociologists and policy makers have long argued that residential segregation by race or income reinforces inequality, in part, through the influence of neighbourhoods on social interactions (e.g., Wilson, 1987; Massey and Denton, 1993; United States Department of Housing and Urban Development, 2016). We know that social interactions have a large causal effect on economic outcomes (e.g., Duflo and Saez, 2003; Bayer, Ross and Topa, 2008; Dahl, Loken and Mogstad, 2014). We understand very little, however, about the role of neighbourhoods in shaping the social environment faced by individuals.

The goal of this paper is to assess the relationship between physical segregation and social segregation by race. I look specifically at the causal influence of a particular aspect of residential segregation: the physical distance it imposes between members of different racial groups. While neighbourhoods may influence social interactions in a variety of ways, the most fundamental of these channels is proximity itself. Social interactions require travel, and this travel acts as a kind of price; moving two individuals into neighbourhoods that are further apart will reduce the frequency of their interactions by increasing the cost of interactions between them. Because residential segregation implies a higher time cost for inter-racial interactions relative to same-race interactions, it presents a barrier to racial integration. The question I ask in this paper is: how important is this effect?

I answer this question through the use of a thought experiment. Taking the existing pattern of residential segregation as given, I simulate the frequency of inter-racial interactions that would occur if physical distance was the *only* factor influencing individuals' interaction decisions. The difference between this frequency and that which would occur in a perfectly integrated world tells us about the magnitude of the barrier posed by residential segregation. If we have a measure of the actual frequency of cross-racial interactions, we can compare the prediction based on distance alone to the actual degree of social segregation to learn about the relative importance of distance compared to other factors in explaining social segregation.[1]

This exercise does not rely on any assumption that residential location is exogenous to racial preferences. We should expect that individuals with stronger preferences for same-race interactions will sort into more segregated cities and neighbourhoods. These individuals will have a higher degree of social segregation than individuals in less segregated neighbourhoods, both because of their preferences and because of the causal effect of distance. The simulation exercises tells us how this behaviour would change if we eliminated the difference in preferences (and all other factors contributing to social segregation), but kept residential location constant. The exercise therefore tells us about how much social segregation can be attributed to the effect of distance alone.[2]

The key parameter I need for my simulation exercise is an estimate of how much individuals dislike travel. I argue that this parameter captures the causal effect of physical distance on social interactions. The argument is based on a discrete choice, transferable utility model of social interaction decisions. The assumptions of the model imply that the frequency of interactions between two people should be related to the joint surplus created by interactions between them. If we take two individuals and move them 1 km further apart, this surplus will fall because they must now jointly travel 1 km further in order to meet. The size of the decline is equal to the disutility of travel. With an estimate of this parameter, I can model the surplus from interactions as depending only on this factor, and use the equilibrium conditions of the model to perform my simulations.

The estimates of the disutility of travel that I use come from the literature on demand for other spatially differentiated goods, such as gas stations (Manuszak and Moul, 2009; Houde, 2012), coffee shops (McManus, 2007), restaurants (Thomadsen, 2005), movie theatres (Davis, 2006), and liquor stores (Seim and Waldfogel, 2013). The assumption of transferable utility in my model allows me to ignore the fact that, unlike travel to gas stations or coffee shops, travel for social interactions may be two-sided. In order to meet, partners must jointly travel at least the distance between their

[1]My results speak to the effect of distance on the subset of social interactions that occur when an individual is otherwise likely to be at home. They do not, for example, include interactions at work. As I describe in my data section, my results are similar if I restrict my data on social interactions to interactions that occur over the weekend.

[2]It is also possible that individuals living in more/less segregated neighbourhoods have different distaste for travel, in addition to different social preferences. I describe how I address this issue in a later section.

homes. So long as they have a way to transfer utility between them, it does not matter how this distance is split between partners. The key assumption I require in using these external estimates is that individuals' dislike of travel is not context specific; in other words, that any differences in the willingness to travel to a gas station or to meet a friend are driven by preferences over the end activity and not over travel itself.

The results of the simulation exercise suggest that if the average white American cared only about travel costs when making interaction decisions, about 9.0% of their interactions would be with black friends. If matching were random within cities, this frequency would be 11.9%. Therefore, the time costs imposed by residential segregation can account for about a 25% reduction in cross-racial interactions relative to the perfectly integrated ideal.

To understand the relative importance of distance compared to other factors in explaining social segregation, we require some measure of the actual frequency of cross-racial interactions. Unfortunately, there are very few large datasets that contain this type of information. Much work on this subject to date relies on data from specialized contexts (high schools or college dorms, for example) that are not appropriate for examining the impact of residential location for more general populations.[3] A second contribution of this paper is to overcome this challenge by presenting a new measure of inter-racial interactions in American cities. This measure is based on a large sample of Flickr photographs, which I run through face detection and race classification software in order to measure the racial breakdown of individual Flickr users' social contacts. Latitude and longitude coordinates on the photographs allow me to link users to cities and neighbourhoods.

My estimates from the Flickr data suggest that white Flickr users see black friends for about 4.2% of their interactions, on average. Adjusting for observable differences between the Flickr sample and the typical white American suggests that this number is substantially lower for white Americans on average, at about 1.8%. This is 10.1 percentage points lower than what we would expect in the perfectly integrated ideal. Accounting for distance eliminates about 30% of this difference. In other words, travel costs alone can explain close to one-third of the social segregation I observe in my data. Another way to interpret this result is to note that if all other factors affecting individuals' choice of interaction partners were eliminated - including racial preferences, segregation in workplaces and schools, and differences in education and other observables that affect the utility of interactions - but

[3]I describe a small number of other papers that examine the relationship between proximity and social interactions, and my contribution relative to these papers, in a later section.

residential segregation remained at its current level, social segregation would still be about one-third of its current level.

Of course, it is possible that the interactions documented in Flickr photographs are not a representative sample of Flickr users' interactions. This will be true if Flickr users are either more or less likely to take pictures of their other-race friends compared to same-race friends. While I cannot rule out this possibility, I can bound my estimates by assuming either a 0% inter-racial interaction rate or an interaction rate equal to the integrated ideal. These bounds suggest that physical distance accounts for between 25-100% of social segregation in the United States. The qualitative conclusion that physical distance presents an important barrier to inter-racial interactions is therefore not dependent on the estimates from my Flickr data.

Note that my results capture only one part of the effect of residential segregation on social interactions. To the extent that neighbourhoods influence interaction decisions in other ways - through schools, for example - my results will underestimate the effect of residential segregation on social segregation. My results nevertheless suggest that the existing degree of segregation plays an important role in limiting black-white interaction. Reducing residential segregation is therefore a necessary condition for substantially reducing social segregation.

In the next section of the paper, I briefly review the earlier literature on the relationship between proximity and social interactions. In section III, I present a model of social interactions that links the causal effect of physical distance to the disutility of travel, and provides the basis for my simulations. In section IV, I describe the data I use to run the simulations, and to measure actual cross-racial interactions in American cities. In section V, I present my results. Section VI concludes.

2 Literature review

There are three other papers that directly examine the relationship between physical distance and the probability of interacting. Empirically, Marmaros and Sacerdote (2006) and Bayer, Hjalmarsson and Pozen (2009) causally identify the effect of proximity on social interactions in college dorms and prisons, respectively, using random or quasi-random assignment to locations within these contexts. These papers show that individuals who live closer together in these types of residential environments are more likely to interact. In both of these cases, however, it is difficult to assess the quantitative relevance of their results in more general settings.

Patacchini, Picard and Zenou (2015) is the only other paper that presents a theoretical model in which the probability of interaction is causally affected by distance. The authors model the returns to interaction as increasing in the partners' interaction rate (which they interpret as social capital), and show that more centrally located agents will interact more often. Using the Add Health dataset (a survey of teenagers that asks individuals to nominate up to 5 friends within the same survey), they show that agents are more likely to be friends when they live closer together, and that more physically central agents appear to be more central in the network as well. While the model I present below is similar in spirit to the model presented in Patacchini, Picard and Zenou (2015) (except for the latter's emphasis on increasing returns to social capital, which I do not consider), my model is better-suited to the kind of quantitative exercise I propose, because it produces equilibrium equations based on a very small number of parameters. Our papers also differ in the outcomes they consider: Patacchini, Picard and Zenou (2015) examine the implications of their model for the level of social capital in a city, and its relationship to transportation costs, while I examine the consequences of residential segregation in predicting social segregation between racial groups. Finally, the data I use to examine the predictions of my model cover a wider segment of the population, and are much larger - the Add Health dataset contains only about 1500 individuals with information on both interactions and residential location.

Another paper that considers a similar research question to mine is Davis et al. (2016), which examines the role of spatial frictions in generating racial segregation in restaurant visits, using data from Yelp. The authors estimate the distaste for travel using the frequency of restaurant visits based on distance from the user's home, work or commute path. They find that spatial frictions account for about half of the observed segregation in restaurant choices. It is difficult to interpret their measure of spatial frictions as reflecting purely the effect of distance, however, because consumers' preferences over restaurants may be correlated with these distance measures; a similar concern in my context motivates my use of an external estimate for the distaste of travel. Additionally, while their paper documents segregation in consumption choices, my paper is able to measure segregation in interactions directly.

Finally, there is a large literature that attempts to assess the importance of physical distance in the market for spatially differentiated goods, such as gas stations, movie theatres, coffee shops and liquor stores (e.g., Thomadsen, 2005; Davis, 2006; McManus, 2007; Manuszak and Moul, 2009; Houde, 2012; Seim and Waldfogel, 2013). The model of social interactions I present below is analagous to the structural models of supply and demand presented in this literature; the key difference is that both the “supplier” and the “customer” travel in my case. I provide a detailed overview of this literature and its results with respect to the disutility of travel in the data section below.

3 Model

In this section, I present a discrete choice, transferable utility model of social interactions adapted from the marriage matching model of Choo and Siow (2006). In this model, each agent resides at a location in a city, and makes a decision each period about whether to interact, and with whom. The transfer that clears the market for interactions in the model is the choice of meeting point, which affects utility because agents are assumed to dislike travel.[4] The transferable utility assumption implies that the equilibrium frequency of interactions between any two groups of agents will depend only on the joint surplus that is created by interactions between them, and on population supplies. All else equal, the joint surplus created by interactions will be declining in the physical distance between two agents, because they must jointly travel this distance in order to meet.[5] The causal effect of distance on social interactions can therefore be captured by agents’ disutility of travel.

In reality, we should expect that agents will have preferences over the observable and unobservable characteristics of their interaction partners. In the version of the model I present here, however, I assume that distance is the *only* factor influencing agents’ utility from interactions. This is because I want to use the model to simulate the pattern of cross-racial interactions based only on physical distance. I show that I can predict the pattern of neighbourhood-by-neighbourhood interactions that would occur in this case, so long as I know agents’ disutility of distance. Using information on the population distribution across neighbourhoods, this can be aggregated into predictions about the frequency with which each individual will interact with people of different races.

[4]The results of the model will be similar so long as there is any way for agents to transfer utility between them (through choice of activity, who pays, etc.)

[5]In the model, I assume that agents always meet somewhere on the line in between them. In real life, agents may choose to visit locations elsewhere in the city. This can be incorporated into the model by allowing agents to have direct preferences over particular locations, which makes this an imperfectly transferable utility model in the manner of Galichon, Kominers and Weber (2016). This extension will not fundamentally change the predictions of the model, however, because the joint surplus of an interaction will still decline in the distance between any two agents’ homes; this is the minimum distance that they must jointly travel.

3.1 Model setup

Agents live in a city, which is represented by a [0,1] line. Each day, agents must decide whether to interact with anyone, and if so, with whom to interact. Because agents are assumed to care only about distance, the only relevant partner characteristic is residential location.

If two individuals decide to interact, they meet at a third location in the city m, which is somewhere on the line in between them. The particular location m that is chosen will adjust to ensure that the market clears.

If an individual agent g who lives at a location l_i interacts with an individual who lives at l_j at a meeting point m, his utility is:

$$U_{gij} = U^s - 2\delta d(l_i, m) + \epsilon_{gij}$$

where U^s is the individual's intrinsic utility from socializing (assumed to be identical for everyone); $d(l_i, m)$ is the physical distance between the agent's home and the meeting point; δ is the disutility of distance; and ϵ_{gij} is an I.I.D. shock with a type I extreme value distribution.

The agent may also choose not to socialize at all, which I denote as choosing partner type "0". In order to capture the fact that individuals still travel when spending time alone, I allow agents to receive shocks to the value of particular locations while alone. Normalizing the intrinsic utility from spending time alone to zero, the utility an agent g living at l_i gets from spending time alone at a location m is

$$U_{gi0} = -2\delta d(l_i, m) + \epsilon_{gi0m}$$

3.2 Equilibrium

Following Choo and Siow (2006), the quasi-demand for interactions with agents living at l_j at meeting point m by agents living at l_i will be:

$$ln(\mu_{ij}^d) = ln(\mu_{i0}) + U^s - \overline{U}_i^0 - 2\delta d(l_i, m) \tag{1}$$

where μ_{ij}^d is the total number of these interactions demanded by agents living at l_i; μ_{i0} is the equilibrium number of agents living at l_i who choose to spend time alone; and $\overline{U}_i^0$ is the average utility that agents living at l_i get when spending time alone (taking the average over m). It is easy to derive an expression for $\overline{U}_i^0$, which will depend only on the agent's location and on δ.[6] In equilibrium, demand for these interactions by agents at l_i must equal the "supply" of these interactions by agents at l_j:

$$ln(\mu_{ij}^s) = ln(\mu_{0j}) + U^s - \overline{U}_j^0 - 2\delta d(l_j, m) \tag{2}$$

The meeting point m will adjust to ensure that this is the case. Setting Equation 1 equal to Equation 2 and solving for the distance travelled by individual i gives:

$$d^*(l_i, m) = \frac{1}{2}d(l_i, l_j) + \frac{1}{4\delta}ln\left(\frac{\mu_{i0}}{\mu_{0j}}\right) + \frac{1}{4\delta}[\overline{U}_j^0 - \overline{U}_i^0] \tag{3}$$

The individual at l_i will travel half of the total distance between the two agents, plus an extra amount reflecting the agent's bargaining power (captured by the latter two terms on the right-hand side of Equation 3.) When there are a large number of unmatched agents living at l_i, relative to the number of unmatched agents living at l_j, the agent at l_i must travel further. This is because there are many close substitutes available to her partner. She must also travel further if the agent at l_j is relatively happier, on average, when spending time alone than she is herself.[7]

Plugging the equilibrium distance equation into the quasi-demand for the agent at l_i gives the equilibrium condition:

$$ln(\mu_{ij}) = \frac{1}{2}ln(\mu_{i0}\mu_{0j}) + U^s - \frac{1}{2}[\overline{U}_i^0 + \overline{U}_j^0] - \delta d(l_i, l_j) \tag{4}$$

Finally, to close the model, note that there is an adding up constraint:

$$\mu_{i0} + \Sigma_k \mu_{ik} = f(l_i) \tag{5}$$

where $f(l_i)$ is the population living at l_i. Rearranging Equation 4 and plugging in the adding up constraint gives:

[6] The exact expression is $\overline{U}_i^0 = log(\int_m e^{-2\delta d(l_i, m)})$. This term captures the fact that, while agents in the center of the city will tend to have a higher equilibrium valuation of social interactions (because they are closer to more people, which minimizes travel costs), they also have a higher valuation of spending time alone.

[7] In a more general case of the model, where agents were not assumed to have identical preferences over all interactions, this term would expand to include differences in the agents' valuation of the interaction as well.

$$ln\left(\frac{\mu_{ij}}{\sqrt{(f(l_i)-\Sigma_k\mu_{ik})(f(l_j)-\Sigma_k\mu_{kj})}}\right)=U^s-\frac{1}{2}[\overline{U}_i^0+\overline{U}_j^0]-\delta d(l_i,l_j) \tag{6}$$

The term on the left hand side of this equation can be thought of as similar to an interaction rate between neighbourhoods i and j; rather than scaling by the total population of these neighbourhoods, however, it is scaled by the number of available partners. The right hand side is equal to the per-partner surplus created by interactions between neighbourhoods i and j, relative to what each partner would get from spending time alone.

If we let N be the number of possible locations (neighbourhoods), this condition gives us a system of $\frac{(1+N)N}{2}$ equations. Denote the right-hand side of this equation as π_{ij}, which is the joint valuation of interactions between agents at l_i and agents at l_j. Choo and Siow (2006) show that, given values for π_{ij} and a vector of population supplies $f(l_i), f(l_j)$ for each neighbourhood, there is a unique vector of social interactions (of size $\frac{(1+N)N}{2}$) that will solve this system of equations.

In the remainder of the paper, I will use Equation 6 to simulate the pattern of neighbourhood-by-neighbourhood interactions that would occur in U.S. cities, using this special case of the model. To do this, I need to estimate the π_{ij}, or joint valuations, for each neighbourhood pair. This requires only information on geographic distance and the disutility of travel.[8] Once I have estimates of π_{ij} and the population supplies, I can predict the frequency of interactions μ_{ij} between any pair of neighbourhoods. I can then aggregate these neighbourhood-by-neighbourhood predictions into overall predictions about the racial breakdown of interactions for individuals living in each area.[9] Because this prediction is formed based only on the causal effect of physical distance, it can tell us about the relative importance of residential segregation versus preferences in explaining social segregation by race.

3.3 Extensions

The model can be extended to the case where agents in different neighbourhoods have different disutilities of travel. We would expect this to be true, because agents will sort into neighbourhoods on the basis of this parameter. For example, agents with high disutility of travel should sort into

[8]Technically, I also need an estimate of U^s, the average utility of interacting. However, when I aggregate the simulation results into predictions about the *relative* frequency of same-race vs cross-race interactions, this term will be eliminated. Therefore, I ignore it when running the simulations.

[9]Under the assumption that individuals do not have intrinsic preferences over interactions with different races, each agent should draw randomly from the population of each neighbourhood. This is how I derive the frequency of interactions by race.

denser, more central neighbourhoods. There are also potential interactions between the disutility of travel and individuals' preferences over interaction partners: conditional on preferences, agents with the highest distaste for travel should sort into the most segregated areas.

Let δ_i be the disutility of travel for agents living in neighbourhood l_i. The equilibrium condition in this case is:

$$ln\left(\frac{\mu_{ij}}{(f(l_i)-\Sigma_k\mu_{ik})^{\frac{\delta_j}{\delta_i+\delta_j}}(f(l_j)-\Sigma_k\mu_{kj})^{\frac{\delta_i}{\delta_i+\delta_j}}}\right)=U^s-[\frac{\delta_j}{\delta_i+\delta_j}\overline{U}_i^0+\frac{\delta_i}{\delta_i+\delta_j}\overline{U}_j^0]-\frac{2\delta_i\delta_j}{\delta_i+\delta_j}d(l_i,l_j) \quad (7)$$

As I explain in more detail in a later section, I will be using my Flickr data to predict a separate disutility of travel for each neighbourhood, on the basis of neighbourhood demographics and location. I can then use Equation 7 to perform my simulations.

A second extension is to the case where agents have non-linear disutility of travel, in line with the results of Davis (2006). This may be the case if, for example, individuals switch modes of travel when travelling longer distances. In this case, utility is not perfectly transferable through meeting location choice. I assume that there is some other technology through which agents can perfectly transfer utility, which takes the form of a transfer τ.[10] Instead of bargaining over meeting point, I assume that agents choose the efficient solution and minimize the joint disutility of travel. This implies that agents meet halfway between their homes. The equilibrium condition in this case is:

$$ln\left(\frac{\mu_{ij}}{\sqrt{(f(l_i)-\Sigma_k\mu_{ik})(f(l_j)-\Sigma_k\mu_{kj})}}\right)=U^s-\frac{1}{2}[\overline{U}_i^0+\overline{U}_j^0]-\delta_1 d(l_i,l_j)-\delta_2 d(l_i,l_j)^2 \quad (8)$$

I simulate this equation using adjusted estimates of δ_1 and δ_2 taken from Davis (2006).[11]

4 Data

To perform my simulation, I need three pieces of information. First, I need to know the population of each neighbourhood, by race. Secondly, I need to know the geographic distance between neighbourhoods within a city. Third, I need an estimate of the disutility of travel. In order to assess

[10]This assumption would not change the equilibrium condition in the linear disutility version of the model.
[11]The adjustment process is described in the data section.

the relative importance of distance in explaining social segregation, I also need estimates of the actual frequency of social interactions by race for individuals, along with information about where the individuals live. In this section, I describe where I get each of these piece of information.

4.1 Population distribution

Information on the population distribution by neighbourhood and geographic distances are available from the U.S. Census Bureau. Throughout the analysis, I will define a neighbourhood as a Census tract. I restrict the analysis to pairs of Census tracts within the same Core-Based Statistical Area (CBSA.)[12] There are 933 CBSAs in the United States (excluding Puerto Rico), of which I use 854 in my main analysis.[13] These CBSAs account for 91.3% of the U.S. population. The mean number of Census tracts one of these CBSAs is 77, ranging from 3 to 4,701.

Table 1 and Figure 1 provide information on the degree of black-white residential segregation in these cities. Table 1 shows the average proportion of the population that is black in the CBSAs, tracts and Census blocks (a unit of measure that is smaller than a tract) inhabited by blacks and whites. There is some segregation apparent even across cities: while the average white person lives in a city that is 11.9% black, the average black person lives in a city that is 20.4% black. These discrepancies become larger at lower levels of geographic disaggregation. At the Census block level, the average white person lives in a block that is 5.5% black, while the average black person lives in a block that is 53.5% black.

An alternative measure of segregation is the Duncan index, which measures the fraction of black or white residents within a city that would have to move to produce an even distribution of racial groups over Census tracts. It is calculated using the following formula:

$$D_c = \Sigma_t | \frac{\text{Black}_{tc}}{\text{Black}_c} - \frac{\text{White}_{tc}}{\text{White}_c} |$$

where Black_{tc} is the number of black individuals living in tract t in city c, and Black_c is the total number of black individuals in the city (and similarly for whites). Figure 1 shows the distribution of the Duncan index across the cities in my sample. The mean Duncan index is 0.506, indicating that about half of all residents in a typical city would have to move to achieve perfect integration.

[12]CBSAs consist of "one or more counties and includes the counties containing the core urban area, as well as any adjacent counties that have a high degree of social and economic integration (as measured by commuting to work) with the urban core." (Census Bureau, 2016.)

[13]I describe the restrictions that lead to my main sample in the section on the disutility of travel below.

This ranges from 0.115 to 0.910 across cities.

4.2 Distance

To measure the geographic distance between pairs of Census tracts, I use shapefiles provided by the U.S. Census Bureau. I calculate the great-circle distance between the central latitude and longitude of each pair of tracts within a CBSA. On average, two randomly selected tracts within the same CBSA are 16.9 km apart; the maximal distance between Census tracts within an average CBSA is 64.9 km.

Table 2 shows the mean distance to an average white and black person within the same CBSA, for white and black individuals separately. Somewhat surprisingly, the average white person lives *closer* to the average black person than they do to other white people. Figure 2 shows how this finding can be reconciled with the substantial degree of residential segregation observed in U.S. cities, using Los Angeles as an example. In Los Angeles, the black population is concentrated in the center of the city. As a result, white individuals in Los Angeles tend to live closer to the average black person (who is likely to live in the city center) than they do to other white people (who are concentrated on the periphery.) A similar pattern holds in many other U.S. cities. Note, however, that this does not imply that distance doesn't play a role in generating social segregation. Although I model the disutility of distance as linear, my simulations produce interaction frequencies that decline non-linearly with the distance between neighbourhoods. It is therefore not the distance to the average black or white person that matters in this context, but the frequency of black and white people in a person's immediate neighbourhood.

4.3 Disutility of travel

A number of papers have estimated individuals' disutility of travel in the context of estimating demand for movie theatres (Davis, 2006; Thomadsen, 2005), liquor stores (Seim and Waldfogel, 2013), coffee shops (McManus, 2007), and gas stations ((Manuszak and Moul, 2009; Houde, 2012). The typical strategy of these papers is to examine how much consumers are willing to pay, in terms of price, to avoid extra travel to a location that is further away. A key assumption for identifying the distaste for travel in this way is that consumers otherwise value the competing locations similarly; that is, that there is no correlation between a location's distance from the consumer and its unobservable characteristics. In some cases, such as for gas stations near a consumer's commute path, this seems reasonable. In other cases where the assumption is more tenuous, a variety of instruments have been used to try and causally identify the effect of distance.

Table 3 summarizes the findings of these papers. The estimated willingness to pay to avoid a minute of travel varies quite substantially in this literature, both in absolute magnitude (ranging from about $0.10-$0.57 per minute in 2002 dollars) and in relation to average hourly wages (with the hourly valuation ranging from 0.5-2.5 times the average hourly wage.) Broadly speaking, however, the results can be grouped into two sets: one set implying a valuation of time at about the average hourly wage (Davis, 2006; McManus, 2007; Manuszak and Moul, 2009), and another set implying a time valuation of about twice the average hourly wage (Thomadsen, 2005; Houde, 2012; Seim and Waldfogel, 2013). I therefore present results based on both of these values.

To extrapolate the estimates in these papers to my sample of cities, I begin by assuming that individuals in each city value their time at either 1 or 2 times the average hourly wage. Hourly wages are available for some metropolitan areas from the Bureau of Labor Statistics. In order to preserve the majority of CBSAs in my sample, however, I instead impute average hourly wages by using information on state-level hourly wages and the ratio of median income in the CBSA to median income in the state. I will use a similar procedure to impute wages at the tract level below.

I convert the dollar valuation of time to utility terms using the estimates from Houde (2012), which are available in both units of measure. This gives me an estimate of disutility per minute, which I then convert into "per kilometre" format by using information on travel speeds from the Google Maps API.[14] I have sufficient information on income and travel speeds to calculate the disutility of travel for 854 CBSAs. The first column of Table 4 summarizes the implied disutility per kilometre across cities. The mean disutility of travel is around 0.417 per kilometre, which corresponds to a dollar valuation of around $0.48 per kilometre in 2010 dollars.

A limitation of this procedure is that variation across cities is imposed by assumption, not by revealed behaviour. We can get some sense of whether the implied distaste for travel actually corresponds to individuals' travel behaviour by using the travel patterns in my Flickr data. As I explain in the next section, the main purpose of my Flickr data is to measure individuals' cross-racial

[14]Specifically, I choose 10 randomly selected pairs of Census blocks within a CBSA and query the API for a driving time between them on a Saturday afternoon at 3 pm. While I use all Flickr photos in calculating my measure of racial segegration in interactions, the results are very similar if I restrict analysis to photos taken on weekends, when people are more likely to be leaving from home as opposed to work.

interactions. Because the photos are geotagged, however, they also provide some information about how individuals move throughout their home cities. Table 5 shows the relationship between my predicted disutility of travel at the city level and the fraction of photographs that are taken within 1, 3, 5 and 10 km of a Flickr user's home location.[15] The city-level regressions in the top panel show that cities with a higher estimated distaste for travel have a higher proportion of photos taken very close to home. A one standard deviation increase in the CBSA-level disutility of distance is associated with a 1.5 percentage point increase in the fraction of photos taken within 1 km of home, a 10% increase from a mean of about 16.1%. The same relationship holds at 3 km, with a slightly larger coefficient implying a 2.0 percentage point increase. The coefficients are also positive at larger distances, although they are less statistically significant.

4.3.1 Tract-level variation

It is possible that the disutility of distance varies not only across cities, but across neighbourhoods within a city. There is no clear theoretical prediction as to how this heterogeneity will affect my results; however, it clearly introduces a source of potential error into my estimates.

As noted in the model section, it is conceptually straightforward to run versions of my simulation that account for this kind of heterogeneity. To implement these simulations, however, I require estimates of how the disutility of travel varies across tracts. One way to do this is to use information on tract level wages and continue to peg the disutility of travel to 1 or 2 times the average hourly wage.[16] The second column of Table 4 provides summary statistics on the tract-level disutility of travel calculated in this way, while the second panel of Table 5 shows how this measure is correlated with Flickr users' travel behaviour. This measure is much more weakly correlated with travel behaviour than the city level measure. The coefficient on the disutility of travel is of the wrong direction for the fraction of photos taken within 1 km of home; while it turns positive for the fraction within 3 and 5 km of home, the size of the coefficient is much smaller than in the city-level regressions.

There are two reasons that the estimates for tract-level disutility of distance may be less reliable than the city-level estimates. First, I have not accounted for differential speeds of travel across Census tracts within a city. This may be problematic, because tracts differ in terms of vehicle ownership

[15] I describe how I infer users' home locations, and provide evidence that I am correctly identifying these locations, in the next section.

[16] Tract level wages are imputed by using the ratio of median income in a tract to median income in the CBSA.

and access to public transit. In particular, the imputation procedure assumes that poorer tracts have a lower disutility of travel, which may not be the case if these individuals have to rely on slower methods of travel. Secondly, even conditioning on income and travel speeds, sorting should induce differences across tracts in the "intrinsic" distaste for travel. This is because individuals who have an unusually high distaste for travel should sort into denser, more central tracts, where travel costs are minimized. Similarly, individuals with a high distaste for travel should live in more segregated areas, even conditional on having the same preferences for own-race interactions. Ignoring this process will cause me to systematically underestimate the disutility of travel for individuals living in dense areas, and overestimate the disutility of travel for individuals living in less dense areas (and similarly for more/less segregated areas.)

As I describe in more detail in the data appendix, I attempt to solve this problem by providing a direct measure of how the distaste for travel varies with demographic characteristics using my Flickr data. As I argue in the appendix, complete travel pattern information (a record of locations visited, and the timing of the visits), along with information about individuals' home and work locations, would make it possible to observe the individuals' distaste for travel directly.[17] The key to the identification strategy is that the travel cost to visit a particular location depends on whether the individual starts from home or work. I can therefore identify individuals' distaste for travel based on whether, conditioning on a particular location's distance from the user's home and from work, the individual is more likely to visit the location when it is relatively "cheap" to do so (after work or on weekends.) See Figure 3 for an example. By aggregating over a large number of individuals, I can examine how this distaste for travel varies with tract demographics.

One problem with this strategy is that I do not have a complete travel record: rather, I have travel information for a set of events that a user deemed worthy of documenting with a photo. This will tend to bias my results, if distance travelled is correlated with the probability of taking a picture (as we might expect if people are more likely to both travel and take a picture on special occasions.) Nonetheless, this exercise may still be informative about the *relative* distaste for travel across neighbourhoods within a city.[18] I therefore peg the mean disutility of travel for each CBSA to be equal to my city-level estimates, and estimate the percentage change in this value associated

[17] As I describe in the next section, I can infer something about Flickr users' home and work locations based on their tendency to take pictures in particular locations at different times of the week.

[18] The key assumption I require for this to be true is that the degree of bias is similar across neighbourhoods once I account for the overall propensity to take photos and post them on Flickr. In particular, I require that any two individuals who post equally frequently on Flickr have similar tendencies to take pictures regardless of their distance from home or work.

with different tract-level characteristics.

Table 6 presents a summary of how the estimated disutility of travel for Flickr users varies with tract characteristics. I regress the predicted disutility of travel for a Flickr user on a measure of tract-level segregation; log median income; and log population density. As expected, individuals living in more segregated areas appear to have significantly higher distaste for travel. There is no significant relationship for density, while individuals living in tracts with higher median income have a significantly higher distaste for travel.

I use the relationships in Table 6 to predict the disutility of travel for all tracts in the 854 cities in my main analysis sample. The third column of Table 4 summarizes the disutility of travel that I infer from this procedure. The last row of Table 5 shows that this measure does a much better job of predicting the proportion of Flickr photos taken close to home than do the measures imputed from average wages. Note that this is not implied by the identification strategy, because the disutility of distance is *not* identified from users' general tendency to travel near or far from home; rather, it is identified from users' tendency to avoid specific locations disproportionately after work, based on whether those locations are accessed more easily from work or home. An increase of one standard deviation in the tract level disutility of travel calculated in this way increases the fraction of photos that an individual takes within 1 km of home by 2.7 percentage points, an increase of 17% relative to the mean in the sample. These relationships remain significant at the 1% level for all distances shown in the table.

4.3.2 Non-linear estimates

A final issue with my estimates of the disutility of travel is that they are assumed to be linear in distance travelled. This may not be the case if, for example, individuals tend to walk short distances and rely on cars or public transit for longer distances. Davis (2006) finds a nonlinear disutility of travel based on movie theatre visits. I construct nonlinear versions of the disutility of distance based on his estimates, scaling his coefficients in line with CBSA-level average hourly wages.[19]

[19] Specifically, his estimates imply that the linear effect of distance ($0.31 in $1996 USD) is equal to about 1.55 times the average hourly wage, using U.S. hourly wages from 1996. The coefficient on the distance squared term is $0.008, equal to about 0.04 times the average hourly wage. These estimates are in miles. I convert them to kilometres, and preserve these ratios when extrapolating these results to other cities.

4.4 Cross-racial interactions

The final piece of information I need to perform the simulation exercise is information on individuals' actual frequency of cross-racial interactions, along with information on where these individuals live. These measures are not available in standard datasets. The publicly available data used in earlier research on social interactions includes the Add Health dataset (a survey of teenagers; e.g., Echenique and Fryer, 2007), the Social Capital Community Benchmark Survey (a survey of individuals living in cities that asks respondents how often they participate in different social activities; e.g., Brueckner and Largey, 2008) and the DDB Needham Lifestyle Survey (a survey that asks similar questions as the SCCBS; e.g., Glaeser and Gottlieb, 2006). Of these, only the Add Health contains information on cross-racial interactions; however, this information is available only for teenagers and has detailed residential information for only a small subsample of respondents.[20]

To measure cross-racial interactions, I instead rely on a novel dataset I have constructed using Flickr photographs. Flickr is a popular photo-sharing website. As of 2013, the site had around 87 million users uploading approximately 3.5 million photos per day (Jeffries, 2013). Flickr users can designate their photographs as "public" or "private"; the company makes a database of all public photographs available to developers. A unique feature of the Flickr database is that all of the metadata attached to the photographs are also accessible. This almost always includes a timestamp, appended by the camera at the time the photo was taken. Additionally, about 5% of photographs have "geotags", which are latitude and longitude coordinates appended by cameras that have access to the internet (smart phones, for example, and higher-end digital cameras.) These geotags allow me to observe individuals' travel patterns and link them to home locations.

4.4.1 Sample construction

I constructed an initial dataset of around 65 million Flickr photographs, all taken within the U.S. between 2006-2015.[21] I require that users are observed primarily in the U.S. As I describe below, I will be attempting to link users to a home location in each year. For this reason, I impose the restriction that users are observed in their most frequently visited CBSA on at least 3 separate days throughout the course of a year. This results in a sample of approximately 170,000 users, who are

[20](Patacchini, Picard and Zenou, 2015) examine the relationship between physical distance and the probability of friendship in the Add Health data, using a sample of about 1500 respondents that have sufficient information on both residential location and social interactions.

[21]I started by pulling a random sample of about 10% of all geotagged photographs taken in the U.S. over this period. Then, I pull every photograph ever taken by the users in this initial sample.

each observed in an average of 1.9 years. This provides a sample of about 325,000 user-years for analysis. These users posted approximately 25 million photos over the sample window.

4.4.2 Home locations

I link users to home locations by assigning them to the modal CBSA in which they take pictures, and to their modal Census tract within the CBSA. I allow users to have a different home location each year. Table 7, Table 8 and Table 9 provide evidence that I have correctly identified users' home locations. Table 7 shows that the home tracts are visited far more often than any other tract. The table shows the number of unique "visits" (day by tract level observations) to the home location and to other Census tracts the user visits. The average user appears in his or her home Census tract on 8.3 days throughout the course of a year; for any other Census tract that the user visits at least once, the mean number of visits is 1.9. For a typical Flickr user, 59.4% of her visits are to the home tract each year; the average among other tracts that she visits is 8.0%.

In Table 8, I show that the surroundings in the home tract are observably different from other tracts the user visits. The table shows the types of venues that appear in the home location and in other visited tracts, using information from the Foursquare database. Foursquare is a service that allows individuals to "check-in" at different locations, providing information to friends and family about where they are. Foursquare maintains a database of venues, which is searchable by latitude and longitude. I search for venues in a 25 metre radius around each photograph, and divide venues into five categories: food and drink (e.g., restaurants, bars, coffee shops), entertainment (e.g., parks, movie theatres, art galleries), stores, offices and residential.[22] I compare the number of venues I find of each type when the user is in his or her assigned home tract and when her or she is elsewhere. The home tract has fewer venues overall than other visited tracts; in particular, it has fewer restaurants, bars and stores. It has more residences and other entertainment facilities, however.

Finally, in Table 9, I use the one piece of information I have on Flickr users - their names - to examine the correlation between the home tract's demographics and the user's demographics. For each user that has a last name on his or her profile (about 55% of the users in my sample), I construct a probability distribution that the user is white, black or other using information on the 1000 most common last names by race in the year 2000 (available from the Census Bureau.)

[22] Foursquare users can add venues to the database; some users add their homes, although this seems to be relatively rare.

Then, I compare this to the fraction of individuals in the user's assigned home tract that are white, black or other. As expected, users with last names that indicate a high likelihood of being white, black or other are assigned to tracts with relatively more of these groups. A one standard deviation increase in the fraction of a user's home tract that is white (about 20 percentage points) increases the probability that a user is actually white by around 1.6 percentage points. The relationship for blacks is smaller, with a one standard deviation increase in the home tract's fraction black (about 15 percentage points) associated with a 0.17 percentage point increase in the probability of being black. For all other groups, a one standard deviation in the home tract percentage (about 14 percentage points) is associated with a 2.5 percentage point increase in the probability that a user belongs to one of those groups.[23]

4.4.3 Social interactions

Once I have linked users to home locations, I measure their social interactions by running their photographs through face detection and race classification software. The face detection algorithm was provided by MIT Information Extraction. Kazemi and Sullivan (2014) report that it has a 95% accuracy rate, with most of the error accounted for by false negatives. The rate of false positives appears to be higher in the Flickr data: based on a sample of 2500 hand-coded photographs, around 8.0% of the faces found by the software are not of human faces.[24] Approximately 12.2% of the photographs in my database (around 3 million) have faces in them. These are posted by around 60,000 users who appear in about 1.7 years on average for a total of about 100,000 user-years. These users account for approximately 2/3 of all photographs in the database. Conditional on a user-year containing any photos with faces, the mean proportion of photos with faces is around 20%.

My measure of interactions is based on the idea that any photograph with faces in it must, in some sense, be documenting a social interaction. As evidence that faces in photographs correspond to actual social behaviour, I compare the frequency of "social" photos (those with any faces in them[25]) to the frequency of social interactions measured at the state-year level in the American Time Use Survey (ATUS). The ATUS asks individuals to keep diaries indicating what they are doing and who they are with at each moment of the day. I use the public-use ATUS file from 2003-2014,

[23] The stronger result for other groups is likely to be due to the fact that last names are more informative for these groups than they are for distinguishing between blacks and whites.

[24] The most common type of photographs that lead to false positives are photographs of statues, art/advertisements, and animal faces. I discuss how I deal with this error in the race classification task below.

[25] The results are similar if I define social photos as those with two or more faces in them.

which contains diary information from about 100,000 individuals. For each individual, I calculate the number of minutes the respondent spent on his or her diary day engaged in "socializing, relaxing and leisure", "eating" or "sports and recreation" with either a non-household family member or a non-household friend.[26] I also construct an indicator for whether a respondent spends non-zero time in these activities on their diary day, and take the mean of both of these measures at the state-year level. Table 10 shows the results from a regression of these variables on the fraction of Flickr photos that are social. Both measures are positively and significantly related to the fraction of social photos, although the relationship is much stronger for the "any social interaction" variable. Moving from the 25th percentile of social photos (around 9.4%) to the 75th percentile (14.0%) is associated with an increase in time spent socializing of 2.6 minutes per day and a 3.1 percentage point increase in the fraction of the population that socializes at all.

To measure cross-racial interactions, I next run all photographs with faces in them through a race classification algorithm. The algorithm itself was provided as part of the Scikit Learn machine learning module for Python. I trained the classifier using the Faces in the Wild database, which is a database of facial photographs designed for studying the problem of unconstrained face recognition.[27] I sorted the photographs from the Faces in the Wild database into three racial groups: black, Asian and other.[28][29] As I discuss in more detail below, I believe that the majority of my Flickr users are white. I therefore measure cross-racial interactions as the fraction of faces in a user's photos that are black.[30]

Table 11 shows the accuracy rates from running the classifier on a subset of the Faces in the Wild database not used for training. To create this table, I used a random sample of about 10% of the photos in the Faces in the Wild database that I excluded for training purposes, and compared the race classifier's predictions about race to my own. The accuracy rate is about 85% for white/other faces, and about 70% for black or Asian faces. While these accuracy rates are much lower than for

[26] The ATUS does not contain information on the race of interaction partners. It also does not contain geographic indicators more detailed than the respondents' state.

[27] The database is available at http://vis-www.cs.umass.edu/lfw/. Unconstrained face recognition involves recognizing faces in contexts that involve non-uniform lighting and poses. I trained the race classifier on this database because its photographs are typically of much higher resolution than the Flickr photographs, which appears to affect performance: algorithms trained on the Flickr data itself have much higher error rates.

[28] Han and Jain (2014) report that there is high inter-subject agreement when using Mechanical Turk workers to sort photographs into age, race and gender groups.

[29] Throughout the remainder of the paper, I refer to faces in the "other" category as "white". This appears to be accurate in the vast majority of cases.

[30] In unreported results, I also examine Asian-white social segregation. Because this appears to be relatively minor, however - whites appear to socialize with Asians at about the rate that would be predicted by CBSA level population frequencies - I focus on black-white interactions throughout the remainder of the paper.

the face detection algorithm, they are in line with the standards in the literature for this type of classification task (see Han and Jain, 2014).)[31]

While the race detection algorithm appears to work reasonably well on the Faces in the Wild database, Table 12 shows that the accuracy rates in the Flickr data are much lower. To produce this table, I hand-coded race for 2500 randomly sampled photographs from the Flickr database with a single face in them, and compared the race classifier's predictions to my own. The algorithm appears to work much less well than it does in the training database, with an overall accuracy rate of about 60% for whites and 50% for blacks and Asians. The lower accuracy of the race-classification algorithm in the Flickr data is likely to be due to the fact that Flickr photographs are typically of much lower resolution than the pictures in Faces in the Wild.

Error in the race classification algorithm is not a problem for my results, because my exercise does not rely on knowing the racial breakdown of faces in any particular photo; rather, I am trying to assess the racial breakdown of faces in a large set of photos. I can use the probabilities in Table 12 to calculate the underlying frequency of black, white and Asian faces in a set of photographs. To understand how I make this calculation, note that the frequency of white faces found by the race detector in my data is about 55%. Based on Table 12, this 55% is comprised of 60.7% of the faces that are actually white; 28.9% of the faces that are actually black; 40.6% of the faces that are actually Asian; and 53.2% of the non-faces (photos of statues, advertisements and animals that are coded as human faces.) If I assume the underlying frequency of non-faces is constant at about 8%, I have an equation with three unknowns - the actual frequency of white, black and Asian faces. By using two additional equations for the frequency of black and Asian faces found by the race detector, I can solve for the value of these variables.[32]

Note that because there is a random element to the coding process (i.e., the error rates are not constant across users), this procedure is unlikely to be accurate for any given user. This is obvious from Table 13, which shows the variation in the fraction of black faces across user-years. The predicted black interaction rate is *negative* for most users, suggesting that the correction process does not work well at an individual level. Over a large number of users, however, it should be accurate. The process suggests that around 5.4% of the faces in my dataset are black.

[31] Accuracy rates are much higher in "constrained" classification tasks, where pose and illumination are constant across subjects.

[32] In practice, I make this adjustment separately by region. The reason for this is that there is a significant variation in both the number of non-faces and in the coding of non-faces across regions, induced by the fact that users in different regions take pictures of different things.

Figure 4 and Figure 5 show how the fraction of black and white faces varies with tract-level demographics. The frequency of white faces is strongly increasing in the proportion of a Flickr user's home tract that is white. Photos taken by users living in strongly non-white tracts have around 40-60% white faces in their photographs; photos taken by users living in tracts that are 100% white have nearly 100% white faces in their photographs. A similar relationship holds for blacks.

The size of the circles in Figure 4 and Figure 5 indicate the number of user-years in each cell. The figures show that the vast majority of Flickr users live in strongly white Census tracts. Furthermore, even those users living in Census tracts with few white individuals have a majority of white faces in their photographs, suggesting that these users are either white or have a high degree of social integration with whites. For this reason, I focus on the behaviour of whites in my analysis. When calculating averages, I weight each user-year by the probability that the user is white based on tract-level demographics. The second column of Table 13 shows how this affects the distribution of black faces across user-years. The weighted mean of black faces among white users is approximately 4.2%.

4.4.4 Sample selection

As noted, Flickr users appear to be disproportionately white. It is also likely that they differ from the typical American along other dimensions. While I do not observe anything about the demographics of Flickr users directly, the results in Table 9 suggest that I can use home tract demographics to proxy for users' unobserved demographic attributes. This provides me with a way to measure the degree of selection in my sample. Table 14 shows the mean of several demographic characteristics in my Flickr users' home Census tracts, compared to the same demographic characteristics of the white U.S. population as a whole. My Flickr users live in larger, denser cities than the typical white American. The cities are similarly segregated; however, Flickr users live in slightly more diverse neighbourhoods, with fewer whites and more Asians. Their tracts have a slightly lower-than-average median age, but higher median income and higher education levels. It will be important to keep in mind, then, that my results from Flickr speak to the behaviour of relatively well-off individuals.

I attempt to correct for sample selection on observables by using the relationships between the fraction of black faces and tract demographics within my Flickr data. Table 15 shows how the fraction of black faces varies with a number of tract-level covariates. The proportion of black faces is increasing in the proportion of black people at both the CBSA and tract level, and is decreasing in the income and education level of the tract. There are significantly fewer black faces in photos taken by users in the Midwest and South Census Divisions, compared to users in the Northeast and West. I can use these relationships to predict the fraction of black faces outside of my Flickr sample, for every tract in the U.S. This exercise suggests that the typical white American has significantly fewer black interactions than do my Flickr users. The predicted black interaction rate for an average white American is just 1.8%.

4.4.5 Photo selection

A second potential issue with the Flickr data is that photos may not be a representative sample of Flickr users' interactions. This will be true if users have a differential propensity to take pictures of black friends relative to white friends. Note that the direction of this potential bias is not clear. Depending on whether a user believes that her social status is heightened by showing inter-racial relationships (if users want to avoid appearing racist, for example) or lowered (e.g., if users worry about racism from others), the proportion of black faces in the photos may be either higher or lower than the true frequency. Without an alternative dataset with which to compare my Flickr results, I cannot rule out this possibility.

Because of this issue, I present lower bound estimates of the impact of distance that are derived from assuming a 0% black-white interaction rate. This is a lower bound estimate because the fraction of social segregation that is "explained" by distance will be smaller when social segregation is assumed to be larger. This set of results do not rely on the Flickr data, and are therefore not subject to this source of bias.

A second issue related to photo selection is that users may be disproportionately likely to take and post pictures of family members - particularly children. Approximately 20% of the faces in the hand-coded photographs are of children. If we assume that these children are of the same race as their parents, than including photos of children may cause me to understate the degree of social integration among friends and other adults. Assuming that 20% of the same-race photographs are of family increases my estimates of the inter-racial interaction rate among friends to 5.3% (for Flickr users) and 2.3% (for a typical white American.) I present versions of my results below that use these figures.

5 Results

In this section, I present the results of my simulation exercise, and compare this to both the random frequency of cross-racial interactions, and the actual frequency of cross-racial interactions from my Flickr data. I first present results using my preferred estimates of the disutility of travel, which are those derived from the Flickr data and pegged to time valuation equal to 1 x the average hourly wage. I then examine how the results change when different estimates of the disutility of travel are used. Table 16 shows the results from my preferred version of the simulation exercise. The first two columns show the results for my sample of Flickr users, while the third and fourth columns show the results for typical white Americans; for each population, I present results using the frequency of black faces in all photos (columns (1) and (3)), and for photos excluding children (columns (2) and (4)). The first row in the table shows the rate of black interaction that I would expect in a perfectly integrated world. This is around 11.1% for Flickr users, and 11.9% for the average white American; these figures represent the average black population frequency in the CBSAs in which each group lives. The third row shows the actual rate of black interaction I estimate from my Flickr data. Focusing on the sample of all faces in columns (1) and (3), this is 4.2% in the Flickr data and 1.8% for the typical whtie American. The gap between the random frequency and the actual frequency - 6.9 percentage points for Flickr users, and 10.1 percentage points for the typical white American - is a measure of social segregation.

The second row of the table shows the frequency of black interaction that I predict from my simulation exercise. This represents the average rate of black interaction that each group would have if its members cared only about distance when making decisions over social interactions. For Flickr users, this figure is around 9.9%. This means that if all other factors affecting the choice of interaction partner - including racial preferences, preferences over characteristics correlated with race, or the effect of schools - were eliminated, social segregation would be about 1.2 percentage points. This represents around 17.4% of the total amount of social segregation for this group.

For the typical white American, the amount of segregation attributable to distance is somewhat higher. The predicted interaction rate is around 9.0%, suggesting that travel costs alone can account for about 2.9 percentage points of social segregation. This is around 29% of the total amount of social segregation observed for this group.

As noted, there is a high degree of error in estimating the actual rate of black-white interaction due to several sources of bias in the Flickr data. The last row of Table 16 therefore provides a lower bound on the effect of distance by showing how large the contribution of distance would be if there was no black-white interaction at all. In this case, the social segregation index for white Americans would be 11.9 percentage points. The 2.9 percentage points accounted for by physical distance make up approximately 25% of this figure.[33] Therefore, the conclusion that physical distance is quantitatively important does not depend on the Flickr estimates.

The second and fourth columns of Table 16 show what happens if I assume that the roughly 20% of faces in my data that are of children are all of the same race as the Flickr user, and exclude these faces from my calculations. In this case, the predicted rate of black interaction is higher, at approximately 5.3% for Flickr users and 2.3% for the typical white American. This causes the estimated degree of social segregation to fall slightly, while the proportion related to distance remains the same; as a result, the part of social segregation attributable to distance rises. Distance is now estimated to account for about 21% of social segregation among Flickr users, and around 30% of social segregation among other white Americans.

Table 17 shows how the results change when I use alternative estimates of the disutility of travel. Column (1) replicates the baseline result from Table 16, for the white population as a whole. Column (2) shows what happens when I ignore the tract-level variation in the disutility of distance, and use the city-level variation only. This makes essentially no difference to the results. Column (3) uses the tract-level variation, but with the mean pegged to 2 times the average hourly wage instead of 1 times the average hourly wage. This causes the simulated black interaction rate to fall from 9.0% to 8.2%. This implies that physical distance explains relatively more (36.6%) of the observed social segregation index. Finally, column (4) uses the non-linear estimate of the disutility of distance from Davis (2006). This again causes the effect of distance to get larger, with a predicted interaction rate of 8.4%. The lower bound estimates (using a zero black-white interaction rate) produced from these simulations range from 24.4% to 31.1%. Therefore, the full range of estimates suggests that the effect of distance explains at least 25%, and potentially up to 37%, of social segregation in the U.S.

Table 18 shows how the results of this exercise vary with city-level segregation. Column (1) repli-

[33]Conversely, if we believed that white and black Americans interacted at the random rate - that there was no social segregation at all - than physical distance would make up more than 100% of the observed degree of social segregation. This would imply that blacks and whites interacted more than we would expect, given the time costs of cross-racial interactions.

cates the main result, while Columns (2)-(5) show the results for cities divided into four quartiles of the Duncan index of segregation. The table shows that social segregation rises with physical segregation. White individuals living in the least segregated cities see black people about 7.1 percentage points less often than would be expected in a perfectly integrated world; in the most segregated cities, this difference is 11.9 percentage points.

The higher degree of social segregation in physically segregated cities is likely to be the combined result of racial preferences (which result in higher levels of physical segregation), along with the causal effect of physical segregation itself. The fifth and sixth rows of the table show that both of these components are indeed higher in more segregated cities. Travel costs can account for 1 percentage point of social segregation in the least segregated cities and 4.5 percentage points in the most segregated cities. The non-explained component of social segregation also rises with physical segregation, but by a smaller amount: it is 6.1 percentage points in the least segregated cities, and 7.4 percentage points in the most segregated cities. As a result, the percentage of social segregation accounted for by distance is higher in more segregated cities. Among the most segregated cities, distance accounts for 37.8% of social segregation.

The preceding paragraph highlights why sorting into residential locations is not a problem for this exercise. It is very likely the case that whites in highly segregated cities have different preferences than other whites - and that these preferences help generate the high levels of segregation we observe. We should therefore expect these individuals to have both a higher segregation index overall (generated by both segregation and preferences) and a higher absolute contribution of distance (generated by time costs alone.) Assessing the *relative* importance of the latter factor is not dependent on any assumption that preferences are similar across more and less segregated areas, because these differences in preferences are built into the raw segregation index. My results suggest that while these preferences do differ across cities with more and less segregation - with the absolute degree of unexplained social segregation rising with physical segregation - the differences in the time costs of social integration are even larger.

6 Conclusion

My results in this paper are based on a framework that links the causal effect of distance to the disutility of travel. Using estimates of this parameter from other contexts, I show that residential segregation appears to be quantitatively important in explaining the tendency of Americans to socialize with members of their own race. Distance appears to explain between one-quarter to one-third of social segregation among white Americans. Even if all other factors contributing to social segregation were eliminated, black-white social interactions would still be 25-30% below the perfectly integrated level, unless residential segregation were also substantially reduced.

It is important to note that my results are likely to substantially understate the impact of neighbourhoods on social segregation. To the extent that neighbourhoods shape interactions in ways beyond the effect of physical distance - through schools, for example, or through the formation of racial preferences - the causal effect of neighbourhoods is likely to account for a much larger portion of social segregation than the 25-30% implied by my estimates. This suggests that it may be very difficult to eliminate social segregation without first targeting the physical segregation between blacks and whites.

7 Figures and Tables

7.1 Figures

Figure 1: Duncan index of black-white segregation across CBSAs

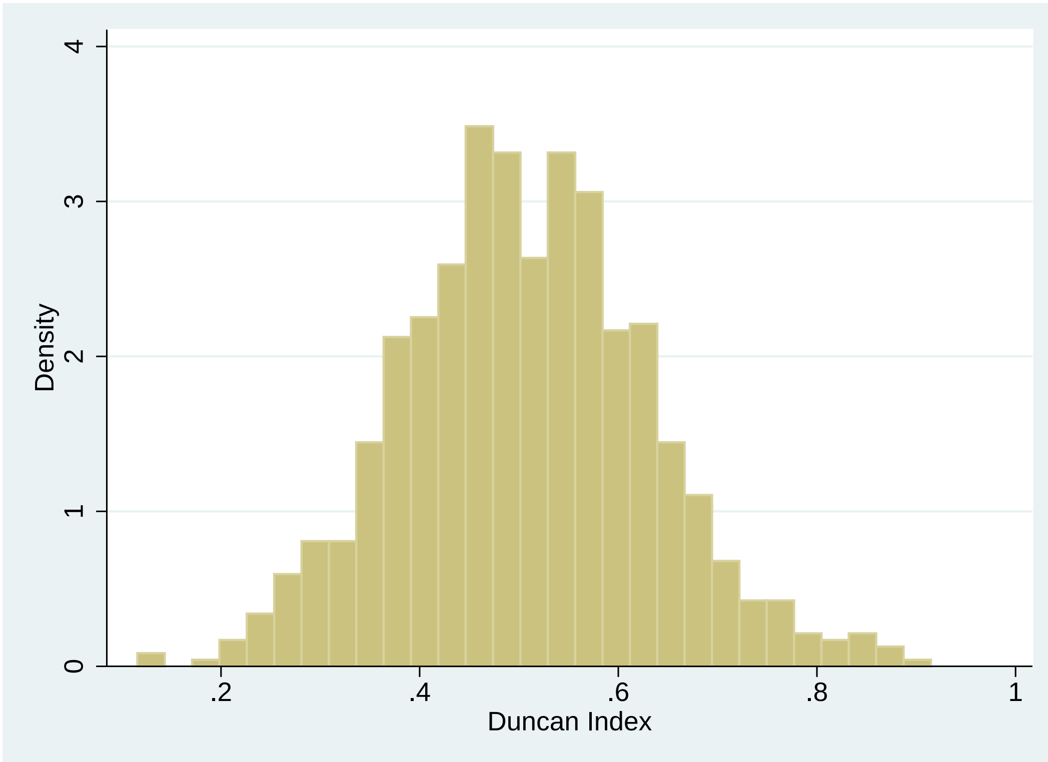

This figure shows a histogram of the Duncan index of black-white segregation across CBSAs in the United States, using information on the population by race of Census tracts from the 2010 Census. See the text for details on the calculation of the Duncan index.

Figure 2: Tract population by race, Los Angeles

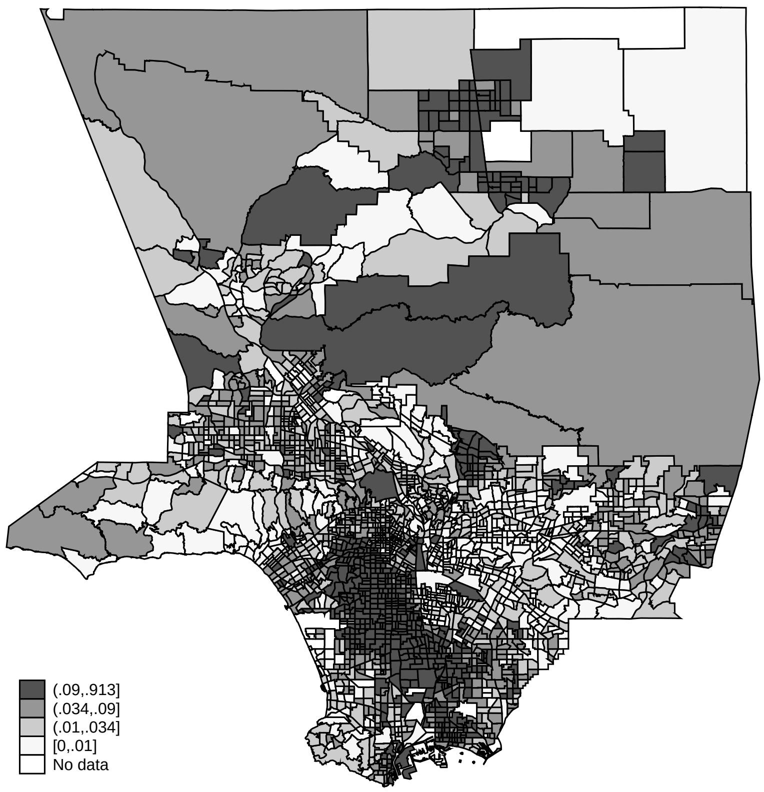

This figure shows the proportion of each Census tract in Los Angeles that is black.

Figure 3: Identifying the disutility of travel: example

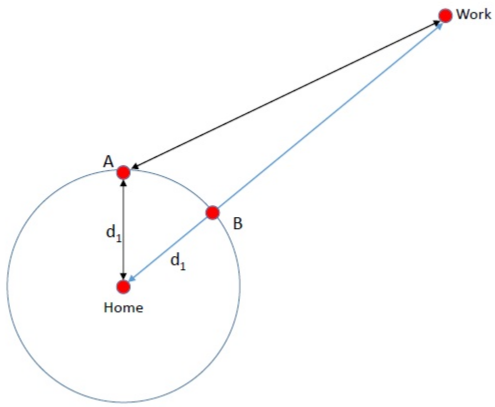

(A)

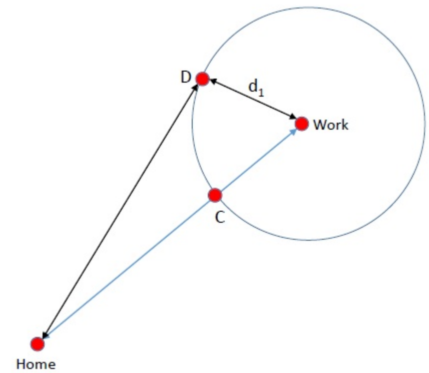

(B)

These figures show how the disutility of travel can be identified using information on travel patterns. In figure A, points A and B are equally costly to the individual on the weekend, when she leaves from home. After work, however, A is more costly than B, because it involves a deviation from her commute path. In figure B, point D is more costly than point C after work; however, its *relative* travel cost is lower after work than on the weekend. This is because visiting D involves a small deviation from the commute path after work, but requires a trip from home and back on the weekend. I can identify the disutility of distance based on individuals' relative probability of visiting points like A and D after work compared to on the weekend. This procedure nets out any correlation between distance from home or work and the individuals' intrinsic taste for the locations.

Figure 4: Neighbourhood racial breakdown and the race of interaction partners

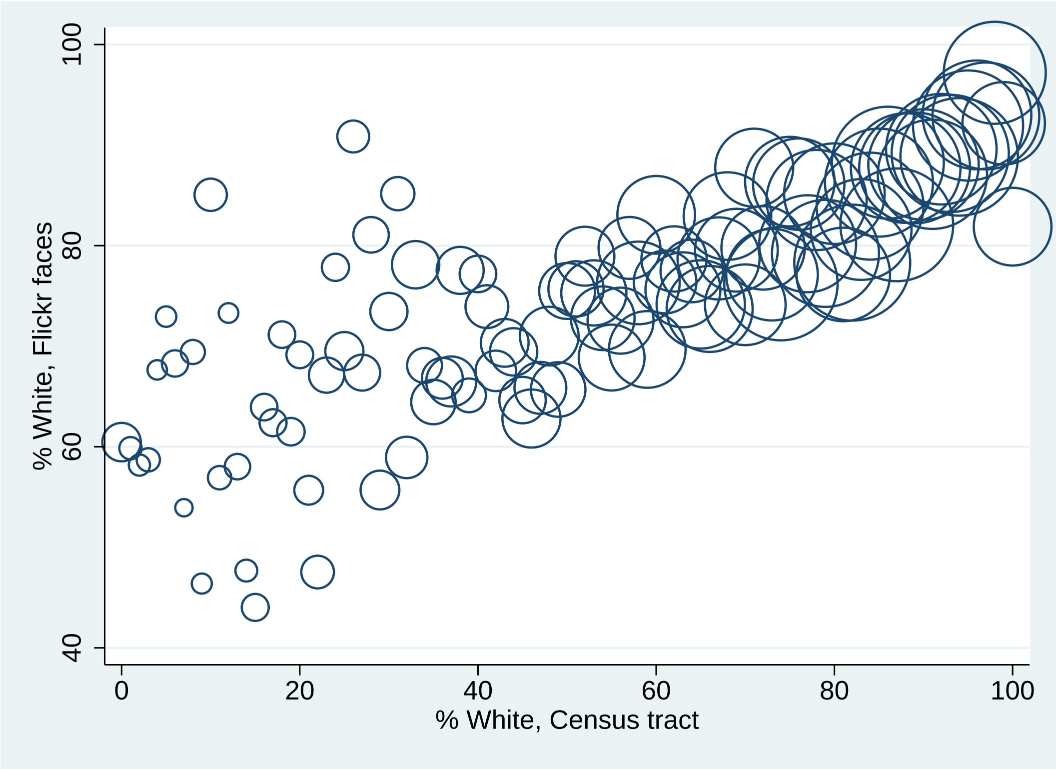

This figure plots the proportion of all faces in a Flickr user's photos that are white against the proportion of her assigned home Census tract that is white. Details on sample selection, the race classification mechanism and the assignment of Flickr users to home Census tracts are available in the Data section. The size of the circles indicates the number of Flickr users in each cell.

Figure 5: Neighbourhood racial breakdown and the race of interaction partners

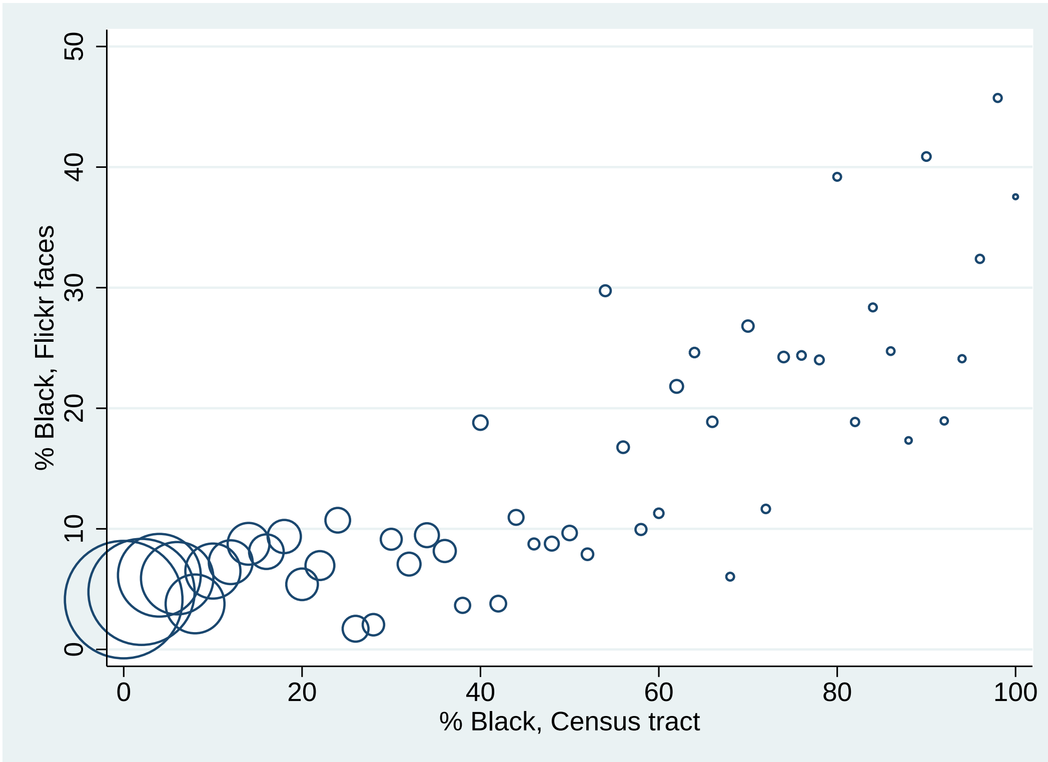

This figure plots the proportion of all faces in a Flickr user's photos that are black against the proportion of her assigned home Census tract that is black. Details on sample selection, the race classification mechanism and the assignment of Flickr users to home Census tracts are available in the Data section. The size of the circles indicates the number of Flickr users in each cell.

7.2 Tables

Table 1: Proportion black at the city, tract and block level, by race

	White	Black
Fraction black - CBSA	11.9%	20.4%
Fraction black - tract	7.3%	46.5%
Fraction black - block	5.5%	53.5%

This table shows the percentage of black people in an average white or black person's CBSA, tract or Census block. The data used to calculate these averages are from the 2010 Census. The sample is the set of all black and white individuals living in one of the 854 cities in my main analysis sample.

Table 2: Distance to the average white/black person, by race

	Distance to the average:	
	White person	Black person
White	28.5 km	27.5 km
Black	27.7 km	23.4 km

This table shows the mean distance to the average white/black person within the same CBSA, for whites and blacks separately. These figures were calculated using great-circle distance between Census tracts, based on shapefiles provided by the U.S. Census Bureau, as well as information on the population of each tract by race from the 2010 Census. The sample is the set of all black and white individuals living in one of the 854 cities in my main analysis sample.

Table 3: Previous estimates of the disutility of distance

	Context	Year(s) of observation	Estimated cost of travel, per minute*	Ratio of travel cost to average hourly wage*
Thomadsen (2005)	Fast food, Santa Clara County	1999	0.49	2
Davis (2006)	Movie theatres, 36 cities	1996	0.23[&]	1
McManus (2007)	Coffee shops, University of Virginia	2000	0.10	0.5-1[#]
Manuszak and Moul (2009)	Gas stations, Chicago & surrounding area	2001	0.18-0.24	0.68-0.91
Houde (2012)	Gas stations, Quebec City	1991-2001	0.10-0.57[@]	0.75-2.50[@]
Seim and Waldfogel (2013)	Liquor stores, Pennsylvania	2005	0.46	1.95

* All dollar estimates are in 2002 USD. Where possible, I use the authors' reported estimates of hourly wages to construct the ratio shown in column (4). Where this is not possible, I use the national hourly wage for the appropriate year, multiplied by the ratio of median income in the relevant geographic area to the median income of the United States.

& Davis (2006) estimates a non-linear function of distance; following Seim and Waldfogel (2013), the reported coefficient is the estimated cost of travelling 3.2 km.

The estimated coefficient is equal to approximately the average wage for students in the relevant geographic market; it is equal to about 0.5 times the average wage for adults in Virginia.

@ The initial estimates reported by (Houde, 2012) are larger than this. His preferred estimates suggest that a time valuation of 4 times the average hourly wage. However, these estimates do not account for traffic. Once I adjust for the average speed of traffic in Quebec City at rush hour (the relevant time, since the estimates examine consumers' willingness to deviate from commute paths), the estimates are reduced to those shown in the table.

Table 4: Disutility of travel: summary

	CBSA variation only	Tract variation	Tract variation - derived from Flickr travel patterns
Mean	0.417	0.471	0.454
Standard deviation	0.104	0.187	0.094
25th percentile	0.355	0.341	0.390
Median	0.402	0.443	0.441
75th percentile	0.464	0.567	0.499
Number of CBSAs	854	854	854
Number of tracts		61,910	61,910

This table shows summary statistics on the estimated disutility of travel, using the three methods described in the text. In the first method, I peg the cost of travel per minute in each CBSA to either 1 or 2 times the average hourly wage, and convert this to a disutility using the estimates in Houde (2012). (Note that only the estimates pegged to 1 times the average hourly wage are shown; the estimates pegged to two times the average hourly wage are twice as high.) The second method is similar, but I extend the variation in the disutility of travel to the tract-level. In the third method, I estimate the distaste for travel directly among Flickr users, and use the relationship between tract-level demographic characteristics and the disutility of travel in this sample to predict the disutility of travel for every tract in my sample. The sample shown in for all measures is the set of cities/tracts for which estimates are available for all measures.

Table 5: Relationship between estimated disutility of travel and travel patterns in Flickr

	Fraction of photos taken within indicated distance of home			
	1 km	3 km	5 km	10 km
Panel 1: City-level, based on hourly wage and travel speeds				
Coefficient	0.144***	0.188***	0.121*	0.112*
	(0.047)	(0.070)	(0.069)	(0.061)
N	854	854	854	854
Panel 2: Tract-level, based on hourly wages and travel speeds				
Coefficient	-0.093***	0.019**	0.030***	0.008
	(0.008)	(0.008)	(0.007)	(0.005)
N	42,402	42,402	42,402	42,402
Panel 3: Tract-level, based on Flickr travel patterns				
Coefficient	0.293***	0.323***	0.281***	0.170***
	(0.016)	(0.016)	(0.014)	(0.011)
N	42,402	42,402	42,402	42,402
Mean of dependent variable:	0.161	0.422	0.563	0.755

This table shows the results from a regression of the mean fraction of photos taken within the indicated distance of Flickr users' homes on the estimated disutility of travel. A increase in the disutility of travel implies that users living within that city or tract dislike travel more. The estimates of the disutility of travel in the first two rows are based on a methodology that scales the cost per minute of travel to 1 times the average hourly wage (imputed from median incomes and state-level hourly wages.) The estimates of the disutility of travel in the last row are based on travel pattern information in the Flickr data, with the mean also scaled to 1 times the average hourly wage and the tract-level variance determined by tract demographic characteristics. The sample in each case is the set of CBSAs or tracts for which there are estimates of the disutility of distance available, and in which at least 1 Flickr user lives.

Table 6: Relationship between disutility of travel and tract characteristics

	Estimated disutility of travel (Based on Flickr travel patterns)
Segregation*	0.260***
	(0.057)
Log density	0.000
	(0.000)
Log median income	0.004***
	(0.000)
Constant	0.043***
	(0.009)
N	191,248
R^2	0.0002

This table shows the results from a regression of Flickr users' estimated disutility of travel on the characteristics of the user's home tract in a given year. A higher disutility of travel implies that the user dislikes travel more. The sample for the regression is the set of user-years in which a user was living in one of the 160 cities for which I have sufficient Flickr data to implement my regressions (for more detail on the regressions, see the Appendix.)
* The measure of segregation I use is the tract's contribution to the Duncan index of segregation: it is the absolute value of the difference between the share of the city's black population that lives within the tract and the share of the city's white population that lives within that tract. Density is the number of individuals per square kilometre.

Table 7: Number of visits to home and other locations in CBSA

	Mean number of visits	Fraction of all visits
Home tract	8.3	59.4%
Other tracts	1.9	8.0%

This table shows the number of unique visits a Flickr user makes to his or her assigned home tract and to other tracts she visits over the course of a year. The sample is a set of approximately 323,000 user-years, which is comprised of approximately 170,000 users who are observed in 1.9 years on average. A user makes a visit to a tract on a given day if she takes at least one photo in that tract on that day. The total number of visits in the sample is approximately 5.7 million.

Table 8: Number of Foursquare venues around photo locations: home tract vs other visited tracts

	Number of venues		
	Home tract	Other visited tracts	Difference
Food & drink	1.203	1.362	-0.160***
			(0.005)
Entertainment	0.619	0.477	0.142***
			(0.003)
Stores	0.543	0.656	-0.113***
			(0.003)
Offices	0.230	0.231	-0.001
			(0.001)
Residential	0.087	0.062	0.025***
			(0.001)
All venues	3.806	4.000	-0.194***
			(0.009)

This table shows the mean number of Foursquare venues within 25 m of a photograph's location, depending on whether that location is within the Flickr user's home Census tract or not. The sample for these calculations is a set of 847,622 photos representing unique visits to Census tracts (i.e. only one photo per day and per tract is kept in the sample), based on a random sample of 1.2 million photographs.

Table 9: Relationship between demographics predicted by last name and home tract demographics

	% probability of being indicated race, based on last name		
	White	Black	Other
% white - home tract	0.078^{***}		
	(0.002)		
% black - home tract		0.011^{***}	
		(0.001)	
% other - home tract			0.177^{***}
			(0.003)
N	175,294	175,294	175,294

The table shows the results from a regression of each Flickr user's probability of being white, black or Asian/other (based on the user's last name) on the percentage of the population in the user's assigned home Census tract that is of the same race. The probabilities based on last name are constructed from a table showing the 1000 most popular last names by race in the year 2000, available from the Census Bureau. To be included in the sample, a Flickr user must have a last name on his or her profile. There are approximately 85,000 users who meet this requirement, which is about half of all users in my sample.

Table 10: Relationship between social photographs and social interactions in ATUS

	Dependent variable:	
	Minutes per day socializing	Fraction of respondents who spend any time socializing
Fraction of Flickr photos that are social	50.1**	0.689***
	(22.1)	(0.114)
N	459	459
R^2	0.011	0.073
Mean of dependent variable	92.0	0.121

This table shows the relationship between measures of social interactions in the American Time Use Survey and the fraction of Flickr photos that are social (contain any faces), at the state-year level. My definition of time spent socializing is the number of minutes engaged in "socializing, relaxing and leisure", "eating" or "sports and recreation" with a non-household family member or friend.

Table 11: Accuracy of race classification algorithm: Faces in the Wild

	Classification:		
	Black	Asian	White
Actual race:			
Black	73.1%	12.9%	14.0%
Asian	3.8%	70.6%	25.6%
White	2.2%	12.7%	85.0%

This table shows the "confusion" matrix for the race classification algorithm for a subset of photos from the Faces in the Wild database not used for training the algorithm. The percentages in the first row show the probability that a black face will be classified as Black, Asian or Other, and similarly for the remaining rows.

Table 12: Accuracy of race classification algorithm: Flickr

	Classification:		
	Black	Asian	White
Actual race:			
Black	45.6%	25.5%	28.9%
Asian	10.6%	48.9 %	40.6%
White	13.5%	25.8%	60.7%
Non-faces	19.3%	27.6%	53.2%

This table shows the classification of faces in my Flickr data, based on the actual race shown (hand-coded) for 2500 photographs. The percentages in the first row show the probability that a black face will be classified as Black, Asian or white, and similarly for the remaining rows.

Table 13: Inter-racial interactions: summary

	Faces, percent black:	
	Unweighted	Weighted
Summary statistics		
Mean	5.4%	4.2%
Standard deviation	69.4%	69.4%
25th percentile	-35.6%	-35.6%
Median	-8.2%	-8.2%
75th percentile	27.1%	26.0%
N	79,722	79,722

This table shows the distribution of the adjusted fraction of black faces in the photographs for each user-year observation in my Flickr data. The first column shows the unweighted distribution, while the second column shows the distribution weighted by the probability that a user is white, based on home tract demographics.

Table 14: Tract demographics: comparison to U.S. population

	White Flickr users	White Americans - all
CBSA population	4,149,674	3,393,587
CBSA segregation index	0.572	0.568
Tract average:		
Density	3,097	1,638
Median age	37.8	38.1
Median income	$33,317	$29,900
% white	79.5	81.4
% black	7.0	7.3
% Asian	6.2	3.9
% Hispanic	12.0	14.4
% No high school	10.2	13.6
% high school	21.5	28.6
% some college	26.0	28.6
% Bachelor's	25.0	18.4
% post-grad	17.3	10.8
N	79,722	198,242,742
Number of tracts	25,064	61,910
Number of cities	807	854

This table shows average home county and tract demographics for user-years in my sample, compared to the averages for the white U.S. population. The averages shown in column (1) are weighted by the probability that a Flickr user is white (based on home tract demographics); the averages in column (2) are weighted by the total white population in each city or tract. Tract demographics are taken from the 2010 U.S. Census, and are shown for tracts in the 854 cities in my main analysis sample.

Table 15: Relationship between tract characteristics and black interaction rate

	Dependent variable:
	Fraction of black faces
% black - tract	0.179^{***}
	(0.028)
% black - CBSA	0.158^{***}
	(0.041)
% high school	-0.211^{***}
	(0.059)
% some college	-0.248^{***}
	(0.045)
% college	-0.046
	(0.048)
% post-graduate	-0.007
	(0.048)
Log median income	-2.310^{***}
	(0.778)
Log density	0.197
	(0.158)
North East	0.193
	(0.833)
Midwest	-9.430^{***}
	(0.795)
South	-6.135^{***}
	(0.850)
N	79,722
R^2	0.007

This table shows the relationship between the percentage of black faces in a Flickr user's photos in a given year and the observable characteristics of his or her assigned home Census tract. All of the tract-level variables are derived from the 2010 Census. The sample is the set of Flickr users that can be linked to home tracts in one of the 854 CBSAs in my main sample, who have photos with faces in them. The regression is weighted by the probability that a Flickr user is white, based on home tract demographics.

Table 16: Simulation results: preferred estimates

	Whites			
	Flickr		Population	
	All	Excluding children	All	Excluding children
	(1)	(2)	(3)	(4)
Black interaction rate				
Random	11.1%	11.1%	11.9%	11.9%
Predicted	9.9%	9.9%	9.0%	9.0%
Actual	4.2%	5.3%	1.8%	2.3%
Social segregation (Random - Actual)	6.9 pp	5.8 pp	10.1 pp	9.6 pp
Explained (Random - Predicted)	1.2 pp	1.2 pp	2.9 pp	2.9 pp
Not explained (Predicted - Actual)	5.7 pp	4.6 pp	7.2 pp	6.7 pp
% explained	17.4%	20.7%	28.7%	30.2%
% explained - lower bound	10.8%	10.8%	24.4%	24.4%
Number of CBSAs	807	807	854	854
Number of tracts	25,064	25,064	61,910	61,910

This table shows the random, predicted, and actual black interaction rate for Flickr users (columns (1) and (2)) and all white Americans (columns (3) and (4)) living in one the 854 CBSAs in my sample. The random frequency is the proportion of the population that is black in an average Flickr user's or white American's CBSA. The predicted black interaction rate is derived from my simulation exercise, and is the frequency of black interactions that would occur if a Flickr user or white American cared only about travel costs when choosing interaction partners. The estimate of the disutility of travel used in this version of the simulation is the tract-level estimate derived from Flickr travel patterns, where the mean time valuation is pegged to 1 x the average hourly wage in a CBSA. The actual frequency of black interactions is derived from Flickr photographs. For Flickr users, this is the average proportion of black faces in white user's photographs; for the white population, it is the adjusted proportion of black faces derived from the regressions in Table 15. Columns (1) and (3) report the frequency of black faces in all photos from this exercise, while columns (2) and (4) report the frequency of black faces in photos excluding children. The lower bound estimates of the impact of distance are derived by assuming no black-white interactions at all.

Table 17: Simulation results: alternative estimates of disutility of travel

	Whites, population			
	Version of δ:			
	Baseline	CBSA variation	2 x wage	Non-linear
	(1)	(2)	(3)	(4)
Black interaction rate				
Random	11.9%	11.9%	11.9%	11.9%
Predicted	9.0%	9.0%	8.2%	8.4%
Actual	1.8%	1.8%	1.8%	1.8%
Social segregation (Random - Actual)	10.1 pp	10.1 pp	10.1 pp	10.1 pp
Explained (Random - Predicted)	2.9 pp	2.9 pp	3.7 pp	3.5 pp
Not explained (Predicted - Actual)	7.2 pp	7.2 pp	7.2 pp	7.2 pp
% explained	28.7%	28.7%	36.6%	34.7%
% explained - lower bound	24.4%	24.4%	31.1%	29.4%
Number of CBSAs	854	854	854	854
Number of tracts	61,910	61,910	61,910	61,910

This table shows the predicted black interaction rate from the simulation exercise for white Americans, using different estimates of the disutility of distance and uses all photos to calculate the frequency of black faces. The baseline estimate is the one shown in Table 16, which relies on the tract-level disutility of distance estimated from Flickr travel patterns. This estimate is pegged to have a mean time valuation of 1 x the average hourly wage. The second row uses only CBSA-level disutility of distance, with the disutility of distance pegged to 1 x the average hourly wage. The third row uses the tract-level Flickr estimates, but pegs the mean valuation to 2 x the average hourly wage. The fourth row uses a non-linear estimate of the disutility of distance, taken from Davis (2006).

Table 18: Simulation results: by segregation quartile

	Whites, population				
	All CBSAs	Segregation quartile I	Segregation quartile II	Segregation quartile III	Segregation quartile IV
	(1)	(2)	(3)	(4)	(5)
Black interaction rate					
Random	11.9%	10.2%	9.0%	10.8%	14.2%
Predicted	9.0%	9.2%	7.7%	8.5%	9.7%
Actual	1.8%	3.1%	1.1%	1.0%	2.3%
Social segregation (Random - Actual)	10.1 pp	7.1 pp	7.9 pp	9.8 pp	11.9 pp
Explained (Random - Predicted)	2.9 pp	1.0 pp	1.3 pp	2.3 pp	4.5 pp
Not explained (Predicted - Actual)	7.2 pp	6.1 pp	6.6 pp	7.5 pp	7.4 pp
% explained	28.7%	14.1%	16.5%	23.5%	37.8%
% explained - lower bound	24.4%	9.8%	14.4%	21.2%	31.6%
Number of CBSAs	854	213	214	213	214
Percentage of white population	100%	11.5%	16.3%	29.1%	43.0%

This table shows the results of my decomposition exercise broken down by the segregation quartile of the CBSA. The

8 References

References

Bayer, Patrick, Randi Hjalmarsson, and David Pozen. 2009. "Building Criminal Capital Behind Bars: Peer Effects in Juvenile Corrections." *Quarterly Journal of Economics*, 124(1): 105–147.

Bayer, Patrick, Stephen L. Ross, and Giorgio Topa. 2008. "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes." *Journal of Political Economy*, 116(6): 1150–1196.

Brueckner, Jan K., and Ann G. Largey. 2008. "Social Interaction and Urban Sprawl." *Journal of Urban Economics*, 64(1): 18–34.

Choo, Eugene, and Aloysius Siow. 2006. "Who Marries Whom and Why." *Journal of Political Economy*, 114(1): 175–201.

Dahl, Gordon B., Katrine V. Loken, and Magne Mogstad. 2014. "Peer Effects in Program Participation." *American Economic Review*, 104(7): 2049–2074.

Davis, Donald R., Jonathan I. Dingel, Joan Monras, and Eduardo Morales. 2016. "How Segregated is Urban Consumption."

Davis, Peter. 2006. "Spatial Competition in Retail Markets: Movie Theaters." *RAND Journal of Economics*, 37(4): 964–982.

Duflo, Esther, and Emmanuel Saez. 2003. "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment." *Quarterly Journal of Economics*, 118(3): 815–842.

Echenique, Federico, and Roland G. Fryer. 2007. "A Measure of Segregation Based on Social Interactions." *Quarterly Journal of Economics*, 122(2): 441–485.

Galichon, Alfred, Scott Duke Kominers, and Simon Weber. 2016. "Costly Concessions: An Empirical Framework for Matching with Imperfectly Transferable Utility."

Glaeser, Edward L., and Joshua D. Gottlieb. 2006. "Urban Resurgence and the Consumer City." *Urban Economics*, 43(3): 1275–1299.

Han, Hu, and Anil K. Jain. 2014. "Age, Gender and Race Estimation from Unconstrained Face Images." *MSU Technical Report*, , (MSU-CSE-14-5).

Hellerstein, Judith, and David Neumark. 2008. "Workplace Segregation in the United States: Race, Ethnicity and Skill." *Review of Economics and Statistics*, 90(3).

Houde, Jean-Francois. 2012. "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *American Economic Review*, 102(5): 2147–2182.

Jeffries, Adrianne. 2013. "The Man Behind Flickr on Making the Service 'Awesome Again'." *The Verge*.

Kazemi, Vahid, and Josephine Sullivan. 2014. "One Millisecond Face Alignment with an Ensemble of Regression Trees." *2014 IEEE Conference on Computer Vision and Pattern Recognition*, 1867–1874.

Manuszak, Mark D., and Charles C. Moul. 2009. "How Far for a Buck? Tax Differences and the Location of Retail Gasoline Activity in Southeast Chicagoland." *Review of Economics and Statistics*, 91(4): 744–765.

Marmaros, David, and Bruce Sacerdote. 2006. "How do Friendships Form?" *Quarterly Journal of Economics*, 121(1): 79–119.

Massey, Douglas S., and Nancy A. Denton. 1993. *American Apartheid: Segregation and the Making of the Underclass.* Cambridge, MA:Harvard University Press.

McFadden, Daniel. 1978. "Modelling the Choice of Residential Location." In *Spatial Interaction Theory and Planning Models.* , ed. Anders Karlqvist, Lars Lundqvist, Folke Snickars and Jorgen W. Weibull, 75–96. Amsterdam, Netherlands:North Holland.

McManus, Brian. 2007. "Nonlinear Pricing in an Oligopoly Market: The Case of Specialty Coffee." *RAND Journal of Economics*, 38(2): 512–532.

Patacchini, Eleonora, Pierre M. Picard, and Yves Zenou. 2015. "Urban Social Structure, Social Capital and Spatial Proximity."

Seim, Katja, and Joel Waldfogel. 2013. "Public Monopoly and Economic Efficiency: Evidence from the Pennsylvania Liquor Control Board's Entry Decisions." *American Economic Review*, 103(2): 831–862.

Thomadsen, Raphael. 2005. "The Effect of Ownership Structure on Prices in Geographically Differentiated Industries." *RAND Journal of Economics*, 36(4): 908–929.

United States Department of Housing and Urban Development. 2016. "Moving to Opportunity for Fair Housing."

Wilson, William Julius. 1987. *The Truly Disadvantaged: The Inner City, the Underclass, and Public Policy.* Chicago, IL:University of Chicago Press.

9 Appendix

9.1 Identifying the disutility of distance from travel patterns

Suppose I had full information on an individual's travel patterns - a complete record of locations visited and the timing of the visits. If I also had information on the individual's home and work locations, I could identify the individual's disutility of distance based on his or her relative probability of visiting particular locations after work, depending on whether that location is "cheaper" starting from work or home. I have some of this information in my Flickr data: I observe visits to particular tracts, and can estimate Flickr users' home and work locations by using their model Census tracts at different times of the week.[34] I use information on their travel patterns on days that they do not appear to socialize - days in which no faces appear in their photographs - to estimate their distaste for travel.

To be more specific, denote an individual i's home location as l_i^h and her work location as l_i^w. Let W_t be an indicator for "after work" (which I will define as weekday evenings in my data.) The distance travelled by an individual to a location m is:

$$D(l_i^h, l_i^w, m, W_t) = 2d(l_i^h, m) * (1 - W_t) + [d(l_i^h, m) + d(l_i^w, m) - d(l_i^h, l_i^w)] * W_t$$

[34] I use the modal Census tract between 9 and 5 pm, Monday-Friday as the user's work location. I eliminate individuals who have "home" and "work" Census tracts that are less than 1 km apart, and impose the restriction that users must be observed in both their home and work locations on at least 3 separate days. This reduces my sample from around 98,000 user-years to about 30,000 user-years.

Using this in a logit equation and denoting $Pr(m,t)$ as the probability that the individual chooses to visit a location m at time t gives:

$$\begin{aligned}Pr(m,t) &= V_i^0(m) - \delta[2d(l_i^h,m)*(1-W_t) + [d(l_i^h,m) + d(l_i^w,m) - d(l_i^h,l_i^w)]*W_t] \\ &= V_i^0(m) - 2\delta d(l_i^h,m) + \delta[d(l_i^w,m) - d(l_i^h,m) - d(l_i^h,l_i^w)]*W_t\end{aligned}$$

where $V_i^0(m)$ denotes the individuals' unobserved valuation of the location m. In general, this may be correlated with distance from either home or work. When estimating this equation, I add both $d(l_i^w,m)$ and $d(l_i^h,l_i^w)$ directly to the model to capture features of the location m or individual i that may be correlated with both $V_i^0(m)$ and distance from home, work and/or the individual's commute. The regression equation then becomes:

$$\begin{aligned}Pr(m,t) = a &+ \beta_h d(l_i^h,m) + \beta_w d(l_i^w,m) + \beta_c d(l_i^h,l_i^w) \\ &+ \delta[d(l_i^w,m) - d(l_i^h,m) - d(l_i^h,l_i^w)]*W_t + \eta_{imt} \quad (9)\end{aligned}$$

The coefficients on the variables $d(l_i^h,m), d(l_i^w,m)$ will not be directly informative, because they will capture both the effect of distance itself and the effect of any correlation between these variables and individuals' utility over particular locations. Instead, the disutility of distance δ is identified from the term in square brackets, which is the "excess distance" that is required to get to a location m after work relative to other times of the week, interacted with an "after work" indicator. To aid intuition, Figure 6 shows an example of a location for which this measure of excess distance is high. In this figure, the travel to points A and B is the same on weekends, because the two points are equidistant from the user's home. However, point A is relatively more costly after work, because it involves a large deviation from the individual's commute path (represented by the blue line between the user's home and work location.) In general, points in the region shaded blue in Figure 7 will be relatively more costly to visit during the week.

It is important to note that δ is not simply identified off of consumers' tendencies to stay on their commute paths, even if this is differentially true on weekday evenings. Figure 8 provides an example. Point C is cheaper than point D after work, because it does not involve a deviation from

Figure 6: High and low excess distance locations (example 1)

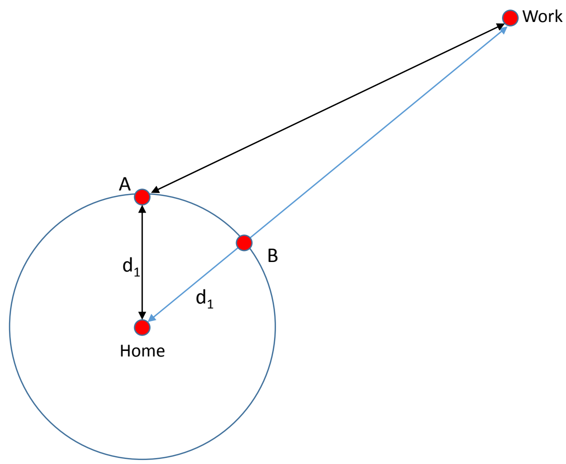

Figure 7: High excess distance region (example 1)

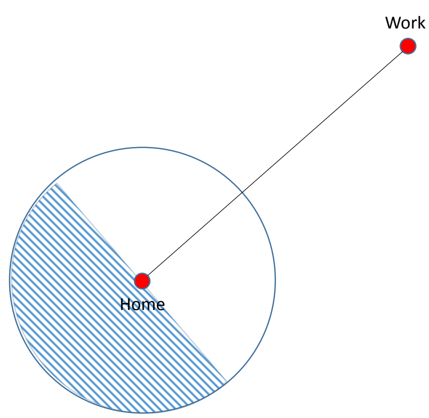

Figure 8: High and low excess distance locations (example 2)

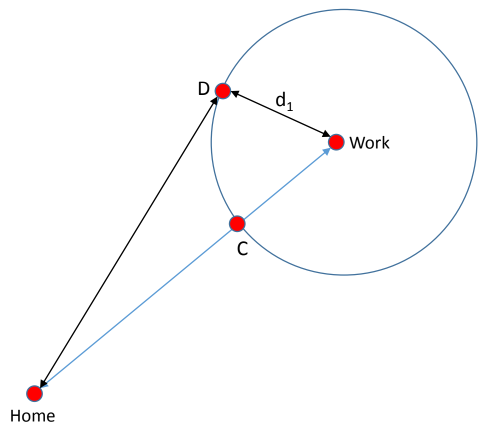

the consumer's commute path. However, it is even cheaper, compared to point D, on the weekends. This is because point D is further from the consumer's home; to visit it on a weekend would involve travelling the entire path from home to D twice. After work, the consumer can visit D with only a minor deviation from her commute. Therefore, the model predicts that the consumer will be more likely to travel to D after work, relative to on the weekends. In general, points in the region shaded blue in Figure 9 will be relatively more costly to visit during the week.

Figure 10 combines the shaded regions from figures 2 and 3 to show the set of regions that have high excess distance, holding both distance from home and distance from work fixed. It is consumers' tendency to disproportionately visit these regions, relative to the non-shaded region *within the same circles* that identifies δ. In particular, note that it is not a problem for identification if the value of regions at a particular distance from work vary depending on the time of the week (if there are happy-hour specials around office buildings, for example), since I am comparing consumers' tendency to visit locations that are equidistant from their offices.

It is possible to allow consumers to have different travel costs depending on the time of week. This is likely to be the case if the speed of travel is lower after work, since this implies that the

Figure 9: High excess distance region (example 2)

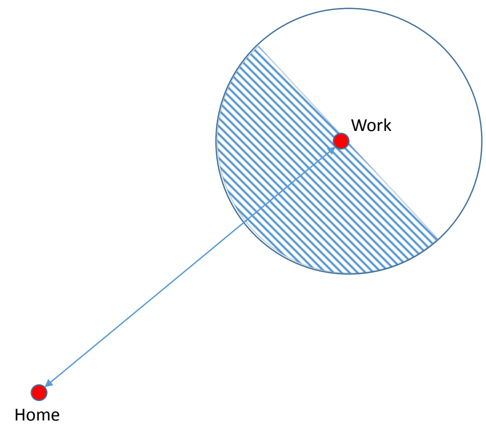

Figure 10: High excess distance regions (combined)

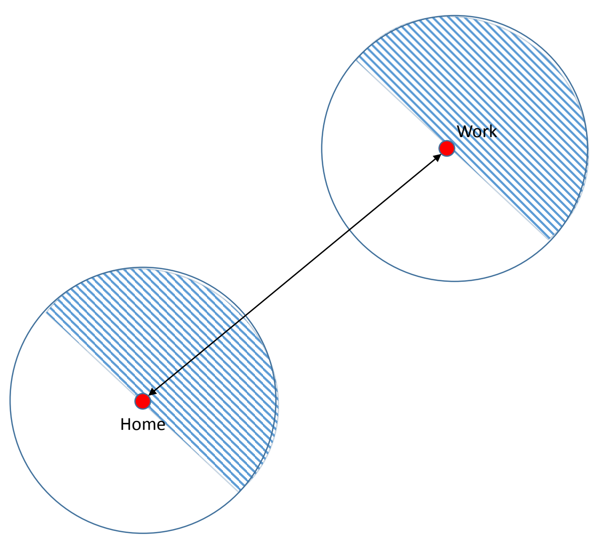

time cost of travelling a particular distance will be higher. If we allow the disutility of travel to be $\delta + \delta_t * W_t$, the estimation equation becomes:

$$Pr(m,t) = \alpha + \beta^h d(l_i^h, m) + \beta^w d(l_i^w, m) + \beta^c d(l_i^h, l_i^w) - (\delta + \delta_t) d(l_i^w, m) W_t \\ - (\delta - \delta_t) d(l_i^h, m) W_t - (\delta + \delta_t) d(l_i^h, l_i^w) W_t \quad (10)$$

By comparing the coefficients on $d(l_i^h, m) * W_t$ and $d(l_i^w, m) * W_t$, we can identify both δ and δ_t.

A problem with using my Flickr data to estimate this equation is that Flickr photographs are not a random sample of travel patterns. In particular, if users are more likely to both travel far and to take a picture on special occasions, it will tend to bias my results towards zero.[35] For this reason, I don't use this methodology to estimate the mean levels of δ; I take these from alternative data sources, as described in the data section. I do, however, use the Flickr data to observe how the disutility of distance varies with observable tract characteristics. I interact the terms identifying the disutility of distance $(d(l_i^h, m) * W_t, d(l_i^w, m) * W_t))$ with measures of observable tract characteristics, and use this to predict a separate disutility of distance for each tract in my sample. I also permit the disutility of distance to vary with the count of Flickr photos produced by each user. The idea is that any differential bias in the estimates of disutility of distance should be captured by individuals' overall propensity to take and post Flickr photos. Among users who post similar numbers of photos overall, I assume that any relationship between tract characteristics and the estimated disutility of distance reflects actual differences in how much the users dislike travel.

Following McFadden (1978) and Davis et al. (2016), I implement the regressions in Equation 9 and Equation 10 in my Flickr data by randomly sampling Census tracts within each user's home CBSA that were not visited by the user. I add these to the set of locations visited by each user to form the user's choice set. I also randomly assign each of these non-visited "observations" a date and an hour. I run the regressions separately by city, allowing the relationship between tract demographics and the disutility of travel to vary individually by city. I restrict the sample to cities that contain at least 10 user-years with at least 10 separate visits in each year, which leaves me with 160 cities.

[35] For example, a user may stay on her commute path most days without ever taking a picture; if she deviates from a path on a single day and takes a picture of it, it will look as though she does not dislike distance.

Once I have tract-level estimates of the distaste for travel, I adjust them so that the mean in each city is equal to the CBSA-level disutility of travel I impute from other sources. Specifically, my estimate of the disutility of travel in tract t in city c is:

$$\delta_{tc} = \overline{\delta_c} + \beta_S * (S_{tc} - \overline{S}_c)) + \beta_D * (D_{tc} - \overline{D}_c) + \beta_N * (N_{tc} - \overline{N}_c)$$

where $\overline{\delta_c}$ is the CBSA mean disutility of distance, produced by adjusting the estimates of Houde (2012) for average hourly wages and travel speeds; S_{tc}, D_{tc} and N_{tc} are measures of tract level segregation, density, and income, respectively; and $\overline{S}_c, \overline{D}_c$ and $\overline{N}_c$ are the city means of these variables. The terms β_S, β_D and β_N are the coefficients reported in Table 6, representing the average relationship between the indicated demographic characteristic and the disutility of travel. Note that while the relationships in Table 6 are based on a sample of users in 160 cities, I use the same relationships to predict the disutility of travel for every tract in each of the 854 cities in my sample.