

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 the impact of physical distance on social segregation by race in the United States. 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 my paper is to assess the importance of this channel in explaining social segregation. I develop a new measure of inter-racial interactions by running a large set of Flickr photographs through face detection and race classification software, allowing me to measure the racial breakdown of individual Flickr users' friends. Geotags on the photos allow me to link Flickr users to neighbourhoods. I use this dataset to show that interactions are highly segregated by race: although Flickr users - who appear to be predominantly white - live in cities that are about 11.6% black, on average, only 5.6% of the faces in my photos are black. I next ask how much of this social segregation can be attributed to the causal effect of distance. Using a transferable utility matching model, I argue that the causal effect of distance on social interactions is captured by consumers' distaste for travel, just as in the case of other spatially differentiated goods. Based on external estimates of this parameter, as well as my own estimates derived from travel patterns in the Flickr data, I use my model to simulate the frequency of cross-racial interactions that would occur if *only* distance mattered in determining individuals' choice of interaction partners. By comparing the level of social segregation that would be predicted purely by physical distance to the actual level of social segregation I observe in my data, I can assess the relative importance of distance in driving individuals' social behaviour. I estimate that 25-45% 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 should be expected to increase social segregation, all else equal. The question I ask in this paper is: how important is this effect?

The first challenge I face in answering this question is a lack of data that contain information on both social interactions by race and on residential location. There are very few large datasets

that measure inter-racial interactions at all. 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.¹ The first 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. This permits to me to directly examine the relationship between residential segregation and social segregation, as in Figure 1. This first panel of this figure plots the fraction of white faces in an individuals' photos against the fraction of their neighbourhood that is white; the second panel presents the same relationship for black faces/black neighbourhoods. As expected, individuals' social contacts reflect the racial breakdown of their neighbourhoods.

Of course, Figure 1 does not imply that these relationships are driven by physical distance. A second challenge I face in assessing the importance of physical distance is identifying the causal effect of distance on social interactions. Clearly, it will be very difficult to infer anything about this parameter by observing individuals' social behaviour. Individuals with stronger preferences for same-race interactions are likely to live in more segregated areas, meaning that it is difficult to attribute any difference in their rate of inter-racial interactions to the effect of distance alone. I argue, however, that it is not necessary to use social behaviour to estimate the causal effect of distance on interactions. Distance matters for social interactions because individuals dislike travel. We know how the same distaste for travel affects consumers' demand for 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). I argue that the same parameter estimated in these papers governs interaction decisions.

My argument is based on a transferable utility matching model of interactions. The assumption of transferable utility 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 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 model also provides me with a way

¹I 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.

to quantitatively assess the importance of physical distance, given an estimate of the distaste for travel. The equilibrium condition from the model says that the frequency of interactions between individuals in any two neighbourhoods will be related to the joint surplus created by interactions between them. This joint surplus will decline with physical distance, at a rate equal to the disutility of travel. I use this condition to predict the pattern of inter-racial interactions that would occur if *only* distance mattered in determining individuals' choice of interaction partners. By comparing the level of social segregation that would be predicted purely by physical distance to the actual level of social segregation I observe in my data, I can assess the relative importance of distance in driving individuals' social behaviour.

To implement this exercise, I use a variety of estimates of the distaste for travel, taken from the literature on spatially differentiated goods. One problem with this procedure is that it ignores any differences in the distaste for travel across individuals. These differences may be systematically related to segregation because individuals with a higher disutility of travel will have a stronger incentive to sort into segregated neighbourhoods. My model is easily extended to permit differences across neighbourhoods in the distaste for travel. I implement simulations using this version of the model by directly estimating individual Flickr users' disutilities of travel, using their travel patterns on days that they do not socialize. My identification strategy relies on the fact that the travel costs to particular locations in a city will be different depending on whether the user leaves from home or work; I can identify their tendency to avoid travel by examining the probability of visiting the same location on weekday evenings compared to the same probability on weekends, depending on whether that location is relatively "cheap" to visit after work.² This permits me to examine how the distaste for travel varies with neighbourhood characteristics, and to predict a separate disutility of travel for every neighbourhood in the U.S.

Figure 2 shows the results of this simulation, using my preferred estimates of the distaste for travel. The first sub-figure plots the predicted frequency of black social interactions for white individuals living in cities against the fraction of the city that is black.³ In a world with no residential segregation, the frequency of black interactions predicted by my model would rise one-for-one with

²As I explain below, this procedure is likely to produce estimates of the disutility of distance that are biased towards zero, because both travel time and the probability of taking a photograph are likely to be higher for special or more valued events. I argue, however, that variation across Flickr users on the basis of tract demographics is still revealing.

³Note that, unlike Figure 1, the variables in both panels of Figure 2 are the same: the rate of black interactions on the Y-axis and the proportion of a city that is black on the X-axis. The difference is in the population used to construct the means of the black interaction rate.

the population frequency of black people in a city; this is shown by the 45 degree line.⁴ Given the actual pattern of residential segregation, the relationship between these two variables falls below this line; the coefficient from a regression of the predicted black interaction rate on the proportion of black people in the city is about 0.8. This means that residential segregation alone predicts that whites will see blacks about 20% less often than would be predicted in a perfectly integrated society. The second sub-figure shows the same relationship for blacks. In this case, the predicted rate of black interactions lies above the 45 degree line; the coefficient from a regression of black interactions on the proportion of a city that is black is 1.26. Residential segregation therefore explains about a 26% increase in black-black interactions from the integrated ideal.

Because Flickr users tend to live in highly dense areas, the effect of physical distance is small for this group: even using relatively high estimates of the disutility of distance suggests that physical segregation explains at most about 7% of the observed social segregation among Flickr users. While I do not directly observe the black interaction rate for white non-Flickr users, I can use the relationship between observable neighbourhood characteristics and the frequency of black interactions within my data to predict a black interaction rate outside of my sample. This procedure suggests that, while the typical white American is likely to have fewer black interactions than my Flickr users (the average rate I predict is around 5.1%), more of this social segregation is explained by physical distance. I estimate that between 25-45% of social segregation among white Americans is due to the effect of residential segregation. This suggests that there is scope to significantly increase cross-racial interactions through policies that lead to desegregation. Of course, the exercise also suggests that other factors - schools, or racial preferences - remain very important in explaining Americans' cross-racial interaction behaviour.

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.

⁴Throughout the paper, I consider interactions within cities only.

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, using random or quasi-random assignment to college dorms and prisons, respectively. 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 disutility 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 analogous 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.⁵ 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.⁶ 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

⁵The 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.)

⁶In 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.

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.

To understand why this prediction is useful, consider a world in which *neither* distance nor preferences mattered in explaining the pattern of social interactions. In such a world, all individuals are identical in terms of their preferences for socializing overall, and all individuals like all other individuals equally. In this world, interactions would be perfectly integrated by race: all individuals would have a representative sample of the population among their friends. The actual pattern of interactions deviates from this ideal because of the causal effect of distance, and because of other factors such as racial preferences. The goal of the analysis is to try and learn something about the relative contribution of physical distance to this gap.

The fact that individuals should be expected to sort into neighbourhoods on the basis of social preferences and the distaste for travel does not affect the validity of this exercise. This type of sorting behaviour would make it difficult to learn about the causal effect of distance simply from observing individuals' social behaviour, because the distance between two people is likely to be correlated with their preferences over interacting with each other. Once we know that the causal effect of distance is captured by the disutility of travel, however, we can use estimates of this parameter from other contexts to produce a "fitted" value, or prediction, about the piece of social behaviour that is explained by distance. By comparing this to an actual measure of social segregation, we can learn about the importance of physical distance in driving social behaviour.

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, who to interact with. 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:

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

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

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

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:

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

where μ_{ij}^d is the total number of these interactions demanded by agents living at l_i ; μ_{i0} is the number of agents living at l_i who choose to spend time alone; and \bar{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 \bar{U}_i^0 , which will depend only on the agent's location and on δ .⁷ In equilibrium, demand for these interactions by agents at l_i must equal the "supply" of these interactions by agents at l_j :

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

⁷The exact expression is $\bar{U}_i^0 = \log(\int_m e^{-2\delta d(l_i, m)})$. 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.

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:

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

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.⁸

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

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

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

$$\mu_{i0} + \sum_k \mu_{ik} = f(l_i) \quad (5)$$

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

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

If we let N be the number of possible locations (neighbourhoods), this condition gives us a system of $N \times N$ equations in $N \times N$ unknowns. 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 pattern of social interactions 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-

⁸In 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.

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.⁹ 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.¹⁰ 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 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. While this is not a problem for evaluating the average contribution of residential segregation to social segregation, it does prevent me from making comparisons across individuals, neighbourhoods or cities in terms of how important physical distance is in explaining their behaviour.

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

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

As I explain in more detail in a later section, I will be using my Flickr data to predict a separate

⁹Technically, 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.

¹⁰Under 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.

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 .¹¹ 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:

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

I simulate this equation using adjusted estimates of δ_1 and δ_2 taken from Davis (2006).¹²

4 Data

I need four pieces of information to perform this decomposition. 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. Finally, I 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 and distance

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).¹³ There are 933 CBSAs in the United States (excluding Puerto Rico.) The mean number

¹¹This assumption would not change the equilibrium condition in the linear disutility version of the model.

¹²The adjustment process is described in the data section.

¹³CBSAs 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)

of Census tracts in a CBSA is 72, and ranges from 3 to 4,696.

Information on the population of Census tracts is taken from the 2010 U.S. Census. Table 1 shows summary statistics on population, density and racial breakdown for all tracts within CBSAs. The first panel of the table shows the distribution of characteristics at the tract level. The typical Census tract has a population of 4,225 individuals, with a range of 0-39,248. The average density (population per square kilometre) is 2,150. The distribution of density is highly left-skewed: the mean is above the 75th percentile of the distribution. On average, the population of Census tracts is 72.5% white, 14.0% black and 4.7% Asian. Panel 2 shows the same information, calculated at the population level. The distribution of characteristics in this panel is similar to the distribution of tract-level characteristics. The advantage of moving to the population level is that I can examine differences in the racial breakdown of tracts for white, black and Asian individuals separately. These measures show that there is a substantial degree of residential segregation by race. The typical white person in the U.S. lives in a Census tract that is 81.5% white, compared to 73.3% for all individuals; the typical black person lives in a Census tract that is 46.7% black, compared to 12.7% for all individuals; and the typical Asian person lives in a Census tract that is 21.3% Asian, compared to 4.9% for all individuals.

Figure 4 shows the distribution of a popular index of segregation, the Duncan index. This index measures how many black residents would have to move to produce an even distribution of racial groups over Census tracts within a city. It is calculated using the following formula:

$$D_c = \sum_t \left| \frac{\text{Black}_{tc}}{\text{Black}_c} - \frac{\text{White}_{tc}}{\text{White}_c} \right|$$

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). The mean Duncan index across CBSAs is 0.501, indicating a substantial degree of residential segregation. This ranges from 0 to 0.910 across U.S. cities.

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 each of the 933 CBSAs. On average, two randomly selected tracts within the same CBSA are 16.9 km apart; the maximal distance between Census tracts within

with the urban core.” (Census Bureau, 2016.)

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 5 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.) Note, however, that this does not imply that interactions will not be segregated. Although I model the disutility of distance as linear, my simulations produce interaction frequencies that decline non-linearly with the distance between neighbourhoods.

4.2 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).

Although these estimates are typically derived from studies of single markets (Quebec City gas

stations, for example), it is possible to extrapolate their findings to other contexts by adjusting for differences in hourly wages and (to convert estimates to a cost per kilometre) speed of traffic. This can be done at the cross-city level and, for wages, at the tract level as well.¹⁴ For my simulations, I need the distaste for travel in terms of utility per km. I convert between utility and dollars using the estimates in Houde (2012). I first adjust his coefficient to imply a time valuation of either 1 or 2 times the average hourly wage in Quebec City, and then multiply this coefficient by the ratio of city or tract level hourly wages to the ratio of Quebec City wages in \$2002 USD. This gives me an estimate of disutility per minute; I convert this into “per kilometre” format by using the Google Maps API to estimate the average speed of traffic in each city.¹⁵ Table 4 summarizes the implied disutility per kilometre across cities and tracts. The mean disutility of travel is around 0.410 per kilometre, which corresponds to a dollar valuation of around \$0.42 per kilometre in 2010 dollars.

A limitation of this procedure is that variation across cities and tracts 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 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- and tract-level and the fraction of photographs that are taken within 1, 3, 5 and 10 km of the user’s home location.¹⁶ The city-level regressions show that cities with a higher 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 further distances, although the coefficients are no longer significant. The relationship between tract-level disutility of travel and individual behaviour is much weaker, however. 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

¹⁴Hourly 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 wage by using information on state-level hourly wages and the ratio of median income in the CBSA to median income in the state. I use a similar procedure to impute hourly wages for Census tracts.

¹⁵Specifically, 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 segregation 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.

¹⁶I describe how I infer user’s home locations, and provide evidence that I am correctly identifying these locations, in the next section.

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 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.¹⁷ 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.) While this is not a problem for the interpretation of my overall results, it does prevent me from comparing the results of my decomposition across cities or neighbourhoods.

As I describe in more detail in the data appendix, I attempt to solve this problem by providing a direct measure of how individuals’ 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 distaste for travel directly.¹⁸ 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 distance from home and from work, the individual is more likely to visit a location that is relatively “cheap” starting from work on weekday evenings. See Figure 6 for an example. The problem, however, 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

¹⁷This is also true at the city-level, although possibly to a lesser extent.

¹⁸As 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.

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. 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.

Table 4 summarizes the disutility of travel that I infer from this procedure. Table 6 shows that, as expected, individuals living in denser and more segregated areas appear to have significantly higher distaste for travel, although these relationships are somewhat small in magnitude. Individuals that live in tracts one standard deviation above the mean of my segregation index have a disutility of travel that is 0.008 higher (a 1.8% increase relative to the mean). Individuals living in tracts that are one standard deviation above the mean in terms of log density have a disutility of travel that is 0.002 higher (a 0.4% increase relative to the mean.) There is no significant relationship between income and the disutility of travel.

The last row of Table 5 shows that the measure of the disutility of travel derived from Flickr travel patterns does a much better job of predicting Flickr users' general tendency to stay 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 this measure increases the fraction of photos that an individual takes within 1 km of home by 2.6 percentage points, an increase of 16% relative to the mean in the sample.

Finally, it is possible that the distaste for travel is non-linear. This would be expected 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

¹⁹Specifically, 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,

In the results section, I will present simulations based on the disutility of travel estimated from all four methods (CBSA-level variation based on average hourly wages; tract-level variation based on average hourly wages; tract-level variation based on estimates from Flickr; non-linear disutility of distance based on Davis (2006)); for each of the first three cases, I peg the mean disutility of travel to either 1 or 2 times the average hourly wage, and present both sets of results. Based on the fact that the Flickr estimates do the best job of predicting individual Flickr user’s travel behaviour, my preferred set of results are those derived from this method. However, I will show that the results are similar regardless of which method I use.

4.3 Cross-racial interactions

The final piece of information I need to perform the decomposition is information on individuals’ actual frequency of cross-racial interactions. In order to compare this to my prediction, I also need 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; this information does not, however, include the frequency of cross-racial interactions.²⁰ Additionally, the Add Health dataset measures the social behaviour of teenagers only, and has detailed residential information for only a small subsample of respondents.²¹

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

and preserve these ratios when extrapolating these results to other cities.

²⁰Add Health asks students to nominate up to 5 friends, but does not ask how frequently individuals see their friends.

²¹(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.

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”, 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.

I constructed an initial dataset of around 65 million Flickr photographs, all taken within the U.S. between 2006-2015.²² 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 require that users are observed in this tract on at least 3 separate days throughout the course of a year, and I allow Flickr users to change locations each year. This results in a final sample of about 50,000 Flickr users who are typically observed in 1 or 2 years, on average. These users posted around 17 million photographs over the sample period.

Table 7, Table 8 and Table 9 provide evidence that I have correctly identified users’ home locations. Table 7 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 the types of venues that appear in the home location, using 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.²³ 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.

²²I 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.

²³Foursquare users can add venues to the database; some users add their homes, although this seems to be relatively rare.

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.) 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 1.7 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 0.25 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 2.9 percentage point increase in the probability that a user belongs to one of those groups.²⁴

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 sample selection. Table 10 shows the mean of several demographic characteristics in my Flickr users’ home Census tracts, compared to the same demographic characteristics of the U.S. population as a whole. My Flickr users live in larger cities and denser areas than the typical American. Their tracts have a higher-than-average median age, median income and proportion white or Asian; the mean education level is also higher. It will be important to keep in mind, then, that my results from Flickr speak to the behaviour of relatively well-off individuals.

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.

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

²⁴The 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.

actual social behaviour, I compare the frequency of “social” photos (those with any faces in them²⁵) 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, 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 11 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 14.4%) to the 75th percentile (20.6%) is associated with an increase in time spent socializing of 3.1 minutes per day and a 4.5 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} Because my Flickr users appear to be primarily white (based on last names, home tract demographics and the high proportion of faces in the database that are white), I measure cross-racial interactions as the fraction of faces in a user’s photos that are black.³⁰

Table 12 shows the accuracy rates from running the classifier on a subset of the Faces in the Wild

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

²⁶I cannot use this measure of interactions for my purposes, because 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.

²⁷The 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.

²⁸Han 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.

²⁹Throughout 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.

³⁰In 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.

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 75% for black or Asian faces. While these accuracy rates are much lower than for 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 13 shows that the accuracy rates in the Flickr data are much lower. To produce this table, I hand-coded race for 250 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 for white faces, and slightly less well for black and Asian faces, than it does in the training database, with an overall accuracy rate of about 55%. 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 serious problem in my context, because my exercise does not rely on knowing the racial breakdown of faces in any particular photo; rather, I am trying to assess the proportion of black faces in a large set of photos. I can use the probabilities in Table 13, along with knowledge of the actual frequency of each type of face in my 250 hand-coded photographs, to calculate the probability that a face is actually black, white or Asian conditional on the algorithm’s prediction of its being black, white or Asian. By appropriately weighting the number of faces of each type that are found by the race classifier in a user’s photographs, I can calculate the actual number of faces of each type. This should be accurate over large samples.

Table 14 summarizes this measure of cross-racial interactions.³² The mean number of black faces in my sample is 5.5%, with a standard deviation of 3.0%. Among Flickr users, the average proportion of black people in their CBSA is 11.6%. This is also the rate of black social interactions that would be observed in a perfectly integrated society. The 6.1 percentage point difference between this ideal and users’ actual behaviour is my measure of social segregation. This is the object that I attempt to decompose in the following section.

³¹Accuracy rates are much higher in “constrained” classification tasks, where pose and illumination are constant across subjects.

³²Information on how this measure varies with city- and tract-level characteristics is shown in Figure 1 and Table 18 in the next section.

5 Results

Table 15 shows the results of my simulation for the entire population of the United States that resides in CBSAs. The typical white person in this group lives in a city that is 11.9% black. In a perfectly integrated city, with no racial preferences, we would therefore predict that the average white person had about 11.9% of their social interactions with black people. Typically, however, cities are not perfectly integrated. Once we account for residential segregation, using estimates of the disutility of distance from Flickr as an estimate of the causal effect of distance, the predicted frequency of black interactions falls to 9.0% (using a version of the estimates that pegs the mean disutility of distance to the average hourly wage) or 8.2% (using a version of the estimates that pegs the mean disutility of distance to 2 times the average hourly wage.) Similarly, the typical black person lives in a CBSA that is 20.5% black. In a perfectly integrated world, this would imply that their social interactions are also 20.5% black. Once we account for residential segregation, however, we see that the predicted fraction of interactions that are black rises substantially: to 36.5% or 41.3%, depending on which version of the estimates of disutility of distance I use.

Columns (3)-(6) of Table 15 show how these results vary by the level of black-white segregation in a city. I divide the CBSAs in my sample into quartiles based on their Duncan index of black-white residential segregation. In the least segregated cities, the level of social segregation explained by physical distance is relatively small: about 1.0% for whites (with a 10.1% predicted interaction rate, compared to an 11.1% random interaction rate) and 5.8% for blacks (with a 24.5% random interaction rate and a 30.3% predicted interaction rate.) By comparison, for the most segregated cities, physical segregation predicts a gap of about 4.5% for whites and 20.8% for blacks.

Table 16 shows that the results are similar when I use alternative estimates of the disutility of distance. In the first two columns of this table, I reproduce the results from the first two columns of Table 15 for comparison. In columns (3) and (4), I use CBSA-level external estimates of the disutility of distance, taken from Houde (2012) and pegged to either 1 or 2 times the average hourly wage in a city (adjusted for traffic speeds.) The estimates are almost identical to my main estimates. In columns (5) and (6), I use the same procedure but allow the disutility of distance to vary with incomes at the tract level. While the estimated impact of distance is slightly smaller for both blacks and whites in this case, the differences are quite small. Distance explains a 23.5% reduction in black interactions for whites in this case, relative to a perfectly integrated ideal (using the estimates

pegged to 1 times the hourly wage); my main estimates suggest a 24.3% reduction. For blacks, segregation explains a 72.7% increase in the frequency of black interactions, compared to 78.0% in my main estimates. Finally, in column (7), I use the non-linear estimates taken from Davis (2006) and adjusted in line with CBSA wages. These results are similar to the main estimates using the the disutility of distance pegged to two times the average hourly wage. In sum, it does not appear that my results are sensitive to the exact estimate of the disutility of travel used for the simulation.

These results suggest that the effect of residential segregation in producing social segregation is quantitatively important. Although we have no measure of the frequency of inter-racial interactions for the general population, we can put a lower bound the importance of residential segregation by thinking about a world in which there are no inter-racial interactions at all. In this case, distance would explain about one-quarter of social segregation for whites, and about 20% of social segregation for blacks, with the remainder due to preferences or other factors. As we increase the actual frequency of cross-racial interactions for these groups, the relative importance of distance would rise and the relative importance of preferences would fall.

While I do not know the frequency of inter-racial interactions for Americans in general, I do have a direct measure of inter-racial interactions for my Flickr users. As noted in the data section, this group is not representative of the general U.S. population in terms of either geographic distribution or racial characteristics. My Flickr users tend to live in disproportionately white Census tracts, and to live in larger, denser cities than the typical American; while these cities tend to be slightly more segregated on average, their higher density means that the consequences of this segregation in terms of explaining social segregation may be relatively limited. Table 17 shows that, in fact, residential segregation explains very little social segregation for this group overall. Using the lower estimate of the disutility of distance produces a *higher* predicted rate of black social interactions than would be produced in a perfectly integrated world; this is because, as noted in the data section, the typical white person actually lives slightly closer to black people than to white people on average.³³ If we increase the average disutility of distance to 2 times the hourly wage rate, the proportion of black interactions falls to below random.³⁴ This effect is still quantitatively small, however. Residential segregation explains a 0.4 percentage point decline in the rate of black interactions for this sample,

³³In unreported results, I show that this is more likely to be true in the cities in which Flickr users are concentrated than for the white population as a whole.

³⁴This is because the predicted rate of interactions falls non-linearly with distance in my model, and the non-linearity becomes more pronounced as the disutility of distance rises.

from 11.6% in a perfectly integrated world to 11.2% once we account for the effect of distance.

The fourth row of Table 17 shows the actual rate of black-white interaction for this sample, estimated from my Flickr data. On average, about 5.5% of the faces in my Flickr photos are black. This is a relatively large deviation from either the 11.6% predicted from random matching, or the 11.2% predicted once we account for residential segregation. Even using the higher estimate of the disutility of distance, segregation explains only about 6.6% of the tendency of Flickr users to predominantly socialize with whites.

The fact that physical distance explains less than 7% of social segregation for Flickr users, but at least 25% of social segregation for whites as a whole is unlikely to be due to differences in the distaste for travel. As shown in the data section, it appears that individuals with a higher distaste for travel tend to sort into denser areas, as we might expect from a model of residential sorting. Because my Flickr users live in denser areas than the typical American, these users are likely to have an unusually *high* distaste for travel, which would suggest that the “explained” portion of social segregation should be higher for this group, all else equal. Instead, it must be the case that Flickr users have relatively strong racial preferences or other factors contributing to social segregation.

To some extent, I may be able to use the relationship between observable covariates among my Flickr users and the “unexplained” portion of social segregation to predict what the frequency of cross-racial social interactions may look like for other groups. Table 18 shows the results of a regression of the proportion of black faces in a Flickr user’s photos on tract-level covariates, controlling for the predicted frequency of black interactions from my simulations. The proportion of black interactions is lower in blacker and more segregated cities, but increasing in the interaction between these two variables; this would be consistent with users in black tracts being more likely to actually be black in highly segregated areas. The frequency of black photos declines in income, but increases in education. Somewhat surprisingly, the fraction of black faces rises with age; this is true whether I use median age or the proportion of individuals in the tract in different age groups. People living in denser area have more photos of black faces, while people living in the Pacific have far more photos of black faces than other regions.

It is important to note that, while these relationships are statistically significant, the observables I use do not predict the fraction of black faces well: the R^2 indicates that they explain only about 2.9% of the variation in the proportion of black faces among my Flickr users. This is potentially due to the high degree of measurement error in my racial classification mechanism. So long as this error

is not correlated with tract characteristics, however, the regression coefficients will still be unbiased. With this caveat in mind, I use the relationships shown in Table 18 to predict the frequency of black interactions for individuals in all tracts in the U.S. The results are shown in Table 19.³⁵ If the relationship between tract observables and the frequency of black interactions among Flickr users holds for the general population, we would expect the average American to have a slightly lower rate of cross-racial interactions than the typical Flickr user, at about 5.0-5.1% (depending on the estimate of disutility of distance that is used as a control.) Combined with the fact that the typical white American lives in a city that is slightly more black than the typical Flickr user (11.9% versus 11.6% for Flickr users), this implies that the total amount of social segregation is slightly higher than in the Flickr sample, at about 6.8 percentage points. Because most white Americans live further away from blacks than my Flickr users do, however, a relatively larger part of this segregation is attributable to distance. For both sets of estimates, the effect of segregation is estimated to be around 42.0%.

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 the causal effect of distance on social interactions is potentially large. In particular, residential segregation appears to be quantitatively important in explaining the tendency of Americans to socialize with members of their own race. At a minimum, physical distance explains about 25% of racial homogeneity in interactions among whites; my estimates from the Flickr data suggest that this may be closer to 45%. Because these estimates represent the causal effect of physical distance, they suggest that there is significant scope to influence inter-racial interaction through policies that reduce residential segregation.

Of course, the fact that residential segregation explains 25-45% of social segregation means that the majority of social segregation remains unexplained. Two obvious factors that could account for the remaining degree of social segregation include schools and racial preferences.

Because school attendance is tightly linked to residential location in the United States, any causal effect of schools on cross-racial interactions will closely mimic a neighbourhood effect. The

³⁵Because most of my Flickr users appear to be white, I show the results of this exercise for whites only.

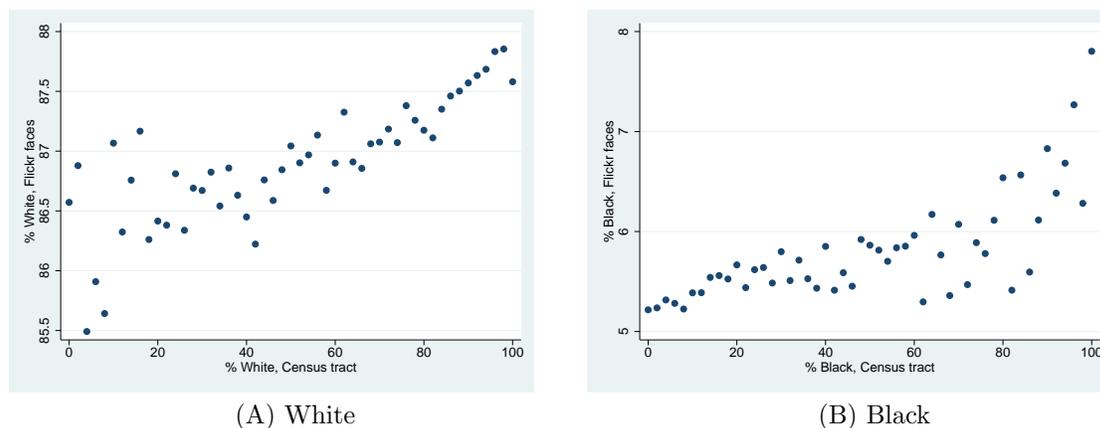
distinction between the two is important, however. The causal effect of physical distance on cross-racial interactions can only be addressed through policies that reduce residential segregation; the causal effect of schools on cross-racial interactions can be addressed through other means.

Finally, it seems likely that racial preferences (or preferences over characteristics correlated with race) explain some degree of the social segregation that remains after accounting for distance. My results provide suggestive evidence that these preferences may be stronger for the demographics represented by Flickr users: relatively white, educated, urban households. To the extent that racial segregation is driven by the location decisions of this demographic (as suggested by the literature on “white flight”), it may prove very difficult to implement policies designed to address residential segregation in a way that does not provoke offsetting responses from this group. On the other hand, recent evidence on the role of exposure in racial preference formation (Carrell, Hoekstra and West, 2015) justifies some optimism. If a short-term increase in interactions can be induced through desegregation policies (which my results suggest is possible), there is reason to believe that this may have even larger effects in the long-term.

7 Figures and Tables

7.1 Figures

Figure 1: 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 (figure (A)) or black (figure (B)) against the proportion of her assigned home Census tract that is white or 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.

Figure 2: The effect of residential segregation on social segregation

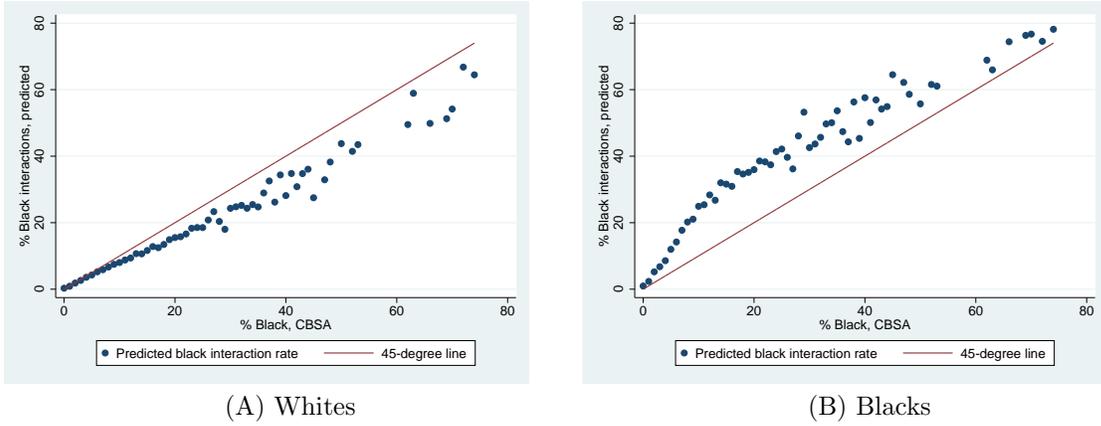
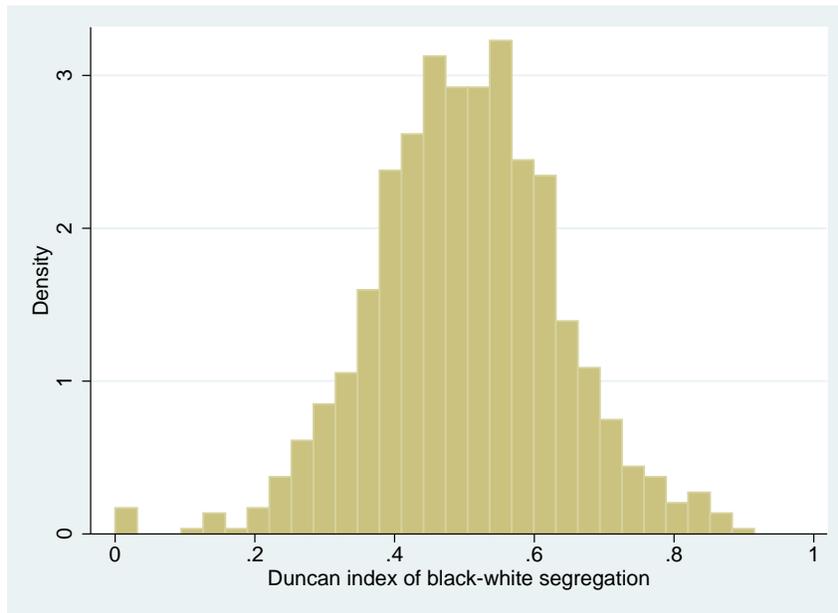


Figure 3: Predicted black interaction rates: white and black Americans

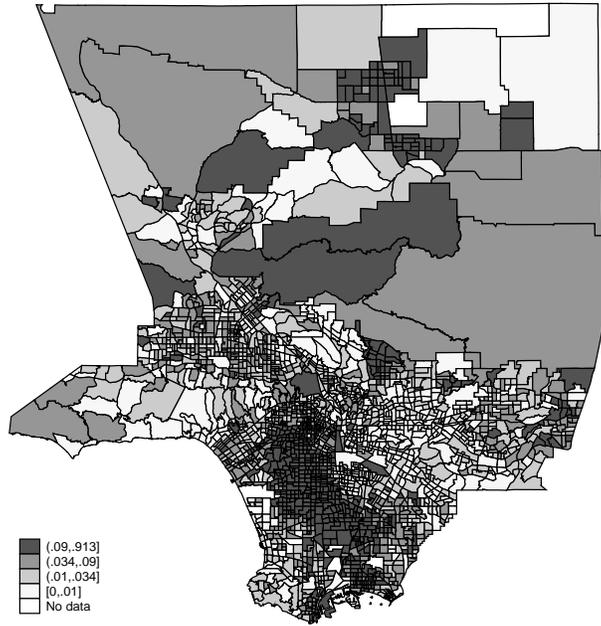
These figures show how the black interaction rate predicted by my simulations varies with the city-level black fraction of the population. The red lines in the graphs are 45-degree lines, which correspond to the rate of black interactions that would be predicted by my model if there was no residential segregation. The blue dots show the actual average predicted rate of black interactions. Details on the simulation procedure are available in the model and data sections.

Figure 4: Duncan index of black-white segregation across CBSAs



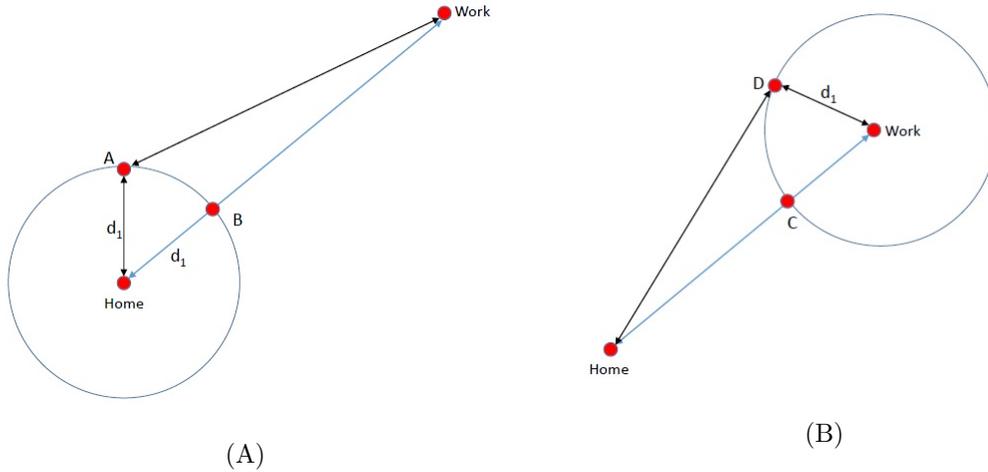
This figure shows a histogram of the Duncan index of black-white segregation across CBSAs in the United States, using information on the population by race of Census tracts from the 2010 Census.

Figure 5: Tract population by race, Los Angeles



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

Figure 6: Identifying the disutility of travel: example



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.

7.2 Tables

Table 1: Population of Census tracts: summary statistics

	Mean	Min	25th percentile	Median	75th percentile	Max
Panel I: tract-level						
Population	4225	0	2899	4020	5321	39248
Density (pop/km ²)	2150	0	205	960	2143	214803
% white	72.5	0.0	60.3	81.5	92.6	100
% black	14.0	0.0	0.8	4.0	15.2	100
% Asian	4.7	0.0	0.1	1.5	5	100
Panel II: population-level						
Population	5100	6	3647	4823	6184	39248
Density (pop/km ²)	2198	0	209	927	2116	214803
% white	73.3	0.0	61.7	81.4	92.2	100
<i>Among whites</i>	81.5	0.1	74.1	86.8	94.1	100
% black	12.7	0.0	0.9	4.0	13.9	100
<i>Among blacks</i>	46.7	0.1	17.4	41.4	77.3	100
% Asian	4.9	0.0	0.3	1.7	5.4	100
<i>Among Asians</i>	21.3	0.1	6.3	14.2	31.2	100

The data for this table are from the 2010 Census. The information is for all Census tracts within one of the 985 CBSAs. There are 67,538 tracts in the CBSAs, which comprised 92.2% of the U.S. population in 2010.

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

	Distance to the average:	
	White person	Black person
White	28.7 km	27.9 km
Black	27.9 km	23.8 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.

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 2002USD. 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	0.418	0.470	0.454
Standard deviation	0.104	0.189	0.095
25th percentile	0.355	0.339	0.390
Median	0.403	0.441	0.441
75th percentile	0.466	0.567	0.499
Number of CBSAs	875	875	875
Number of tracts		64,595	64,595

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.

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
<i>City-level, based on hourly wage and travel speeds</i>				
Coefficient	0.143***	0.173**	0.116	0.094
N	850	850	850	850
<i>Tract-level, based on hourly wages and travel speeds</i>				
Coefficient	-0.093***	0.017**	0.024***	0.000
N	42,921	42,921	42,921	42,921
<i>Tract-level, based on Flickr travel patterns</i>				
Coefficient	0.276***	0.290***	0.248***	0.150***
N	44,093	44,093	44,093	44,093
Mean of dependent variable:	0.161	0.422	0.563	0.755

This table shows the results from a regression of the mean fraction 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.

Table 6: Relationship between disutility of travel and tract characteristics

	Estimated disutility of travel (Based on Flickr travel patterns)
Segregation*	0.407*** (0.057)
Log density	0.001*** (0.000)
Log median income	0.001 (0.001)
Constant	0.005 (0.007)
N	223,148
R^2	0.0004

This table shows the results from a regression of individuals' estimated disutility of travel on the characteristics of the user's home tract. A higher disutility of travel implies that the user dislikes travel more.

* 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.

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 847,622 unique visits to Census tracts (where a visit is a day by Census tract observation), 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	0.085*** (0.004)		
% black - home tract		0.016*** (0.002)	
% other - home tract			0.196*** (0.006)
N	50,258	50,258	50,258

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.

Table 10: Tract demographics: comparison to U.S. population

	Average - Sample	Average - U.S. population
CBSA population	4,972,186	1,080,838
Tract average:		
Density	3,265	2,198
Median age	37.9	37.1
Median income	\$33,948	\$29,046
% white	74.3	73.2
% black	9.3	12.7
% Asian	8.1	4.9
% Hispanic	12.4	16.4
% No high school	11.2	15.2
% high school	21.0	28.3
% some college	25.1	28.2
% Bachelor's	24.5	17.8
% post-grad	18.2	10.5

This table shows county and tract demographics for the home locations of my sample, compared to the averages for the U.S. population. Tract demographics are taken from the 2010 U.S. Census.

Table 11: 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	50.093*** (14.980)	0.773*** (0.075)
N	459	459
R^2	0.024	0.187
Mean of dependent variable	92.621	0.484

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 12: Accuracy of race classification algorithm: Faces in the Wild

	Classification:		
	Black	Asian	Other
Actual race:			
Black	73.1%	12.9%	14.0%
Asian	3.8%	70.6%	25.6%
Other	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 13: Accuracy of race classification algorithm: Faces in the Wild

	Classification:		
	Black	Asian	Other
Actual race:			
Black	50.0%	21.4%	28.6%
Asian	9.5%	61.9%	28.6%
Other	14.4%	30.2%	55.4%

This table shows the classification of faces in my Flickr data, based on the actual race shown (hand-coded) for 250 photographs. 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 14: Inter-racial interactions: summary

	Faces - percent black
Summary statistics	
Mean	5.5%
Standard deviation	3.0%
25th percentile	3.3%
Median	4.6%
75th percentile	6.4%
Population frequency	11.6%
Social segregation (mean)	6.1%
N	83,116

This table shows the fraction of black faces in the photographs for each user-year observation in my Flickr data.

Table 15: Effect of segregation on cross-racial interactions: population

	(1)	(2)	(3)	(4)	(5)	(6)
Black interaction rate						
<i>Whites</i>						
Random	11.9%	11.9%	11.1%	8.7%	10.8%	14.2%
Predicted	9.0%	8.2%	10.1%	7.5%	8.5%	9.7%
Difference (Random - Predicted)	2.9%	3.7%	1.0%	1.2%	2.3%	4.5%
% difference	-24.3%	-31.0%	-9.0%	-13.8%	-21.3%	-31.7%
<i>Blacks</i>						
Random	20.5%	20.5%	24.5%	18.8%	21.5%	19.6%
Predicted	36.5%	41.3%	30.3%	27.7%	35.3%	40.4%
Difference (Random - Predicted)	16.0%	20.8%	5.8%	8.9%	13.8%	20.8%
% difference	78.0%	101.4%	23.7%	47.3%	64.2%	106.1%
Population	All	All	Segregation quartile I	Segregation quartile II	Segregation quartile III	Segregation quartile IV
δ pegged to:	1 \times wage	2 \times wage	1 \times wage	1 \times wage	1 \times wage	1 \times wage

This table shows the frequency of black interactions predicted by a random matching model and by my simulations, for the entire U.S. population living in the 853 CBSAs I use for my simulations.

Table 16: Effect of segregation on cross-racial interactions: population, alternative estimates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Black interaction rate							
<i>Whites</i>							
Random	11.9%	11.9%	11.9%	11.9%	11.9%	11.9%	11.9%
Predicted	9.0%	8.2%	9.0%	8.2%	9.1%	8.4%	8.4%
Difference	2.9%	3.7%	2.9%	3.7%	2.8%	3.5%	3.5%
% difference	-24.3%	-31.0%	-24.3%	-31.0%	-23.5%	-29.4%	-29.4%
<i>Blacks</i>							
Random	20.5%	20.5%	20.5%	20.5%	20.5%	20.5%	20.5%
Predicted	36.5%	41.3%	36.4%	41.2%	35.4%	40.2%	39.8%
Diference	16.0%	20.8%	15.9%	20.7%	14.9%	19.7%	19.3%
% difference	78.0%	101.4%	77.6%	100.9%	72.7%	96.1%	94.1%
Population	All	All	All	All	All	All	All
δ pegged to:	1 \times wage	2 \times wage	1 \times wage	2 \times wage	1 \times wage	2 \times wage	N/A
δ variation:	Tract	Tract	CBSA	CBSA	Tract	Tract	CBSA
Source of δ :	Flickr	Flickr	External	External	External	External	External
Functional form	Linear	Linear	Linear	Linear	Linear	Linear	Non-linear

This table shows how the results of my simulations differ when I use alternative estimates of the disutility of distance. The first two columns correspond to the estimates shown in Table 15. For these columns, the mean level of the disutility of distance is pegged to either 1 or 2 times the average hourly wage in the city (converting to utility terms using the estimates shown in (Houde, 2012)), with tract-level variation estimated from the Flickr data. In the second two columns, I use the estimates from (Houde, 2012) directly, adjusting at the CBSA level for speed of traffic and pegging the mean to either 1 or 2 times the average hourly wage. In column (5) and (6), I use the same procedure, but introduce tract level variation based on median incomes. Finally, in the last column, I use a non-linear estimate of the disutility of distance from (Davis, 2006), adjusting both the intercept and coefficient to maintain the same ratio with average hourly wages as in his paper.

Table 17: Effect of segregation on cross-racial interactions: Flickr users

	(1)	(2)
Random	11.6%	11.6%
Predicted	11.8%	11.2%
% difference	1.5%	-4.0%
Actual	5.5%	5.5%
Social segregation (Random - Actual)	6.1%	6.1%
Explained (Random - Predicted)	-0.2%	0.4%
Unexplained (Predicted - Actual)	6.3%	5.7%
% unexplained	103.3%	93.4%
δ pegged to:	1 \times wage	2 \times wage
δ variation:	Flickr	Flickr
Source of δ :	Tract	Tract
N	83,116	83,116

In this table, I show the results of the decomposition for Flickr users. The results in this table are based on the estimates of the disutility of distance derived from the Flickr travel patterns data, with the mean disutility of distance pegged to either 1 or 2 times the average hourly wage.

Table 18: Relationship between tract characteristics and unexplained black interactions

	Dependent variable:	
	Fraction of black faces	
% black	-0.013*** (0.003)	-0.018*** (0.003)
% white	-0.000 (0.002)	-0.000 (0.002)
CBSA segregation index	0.003 (0.003)	0.003 (0.003)
CBSA segregation index × % black	0.021*** (0.004)	0.023*** (0.004)
CBSA segregation index × % white	-0.006* (0.004)	-0.006*** (0.003)
Log median income	-0.001*** (0.000)	-0.001*** (0.000)
Log mean years of education	0.010*** (0.001)	0.010*** (0.001)
Log median age	0.002*** (0.000)	0.002*** (0.000)
Log density	0.000*** (0.000)	0.000*** (0.000)
Black interactions - predicted	0.014*** (0.001)	0.016*** (0.003)
Region fixed effects	X	X
N	83,116	83,116
R ²	0.029	0.029

This table shows the relationship between the proportion of black faces in a Flickr user's photos and the observable characteristics of his or her assigned home Census tract. I control for the predicted frequency of black interactions from my simulations, so that the regressions can be interpreted as describing patterns in the unexplained frequency of black interactions. In the first column, I construct this control using estimates of the disutility of distance pegged to 1 times the average hourly wage; in the second column, I use estimates of the disutility of distance pegged to 2 times the average hourly wage.

Table 19: Decomposition results: entire white population

	(1)	(2)
<i>Whites</i>		
Random	11.9%	11.9%
Predicted	9.0%	9.0%
Actual (Estimated)	5.1%	5.0%
Social segregation (Random - Actual)	6.8%	6.9%
Explained (Random - Predicted)	2.9%	2.9%
Unexplained (Predicted - Actual)	3.9%	4.0%
% explained	42.6%	42.0%

In this table, I show the results of the decomposition for all whites living in CBSAs in the United States. Estimates of the actual frequency of inter-racial interactions are derived from the relationship between this measure and observable tract characteristics in the Flickr data.

8 References

References

- Bayer, Patrick, Randi Hjalmarrsson, 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.
- Carrell, Scott E., Mark Hoekstra, and James E. West.** 2015. "The Impact of Intergroup Contact on Racial Attitudes and Revealed Preferences." *NBER Working Paper*, No. 20940.

- 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.
- Dufo, 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, Adrienne.** 2013. “The Man Behind Flickr on Making the Service ‘Awesome Again’.” *The Verge*.

- Kazemi, Wahid, 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 Data 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 can 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.³⁶ 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:

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

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

³⁶I 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 83,000 user-years to about 30,000 user-years.

commute. The regression equation then becomes:

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

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 7 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 8 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 9 provides an example. Point C is cheaper than point D after work, because it does not involve a deviation from 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 10 will be relatively more costly to visit during the week.

Figure 11 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

Figure 7: High and low excess distance locations (example 1)

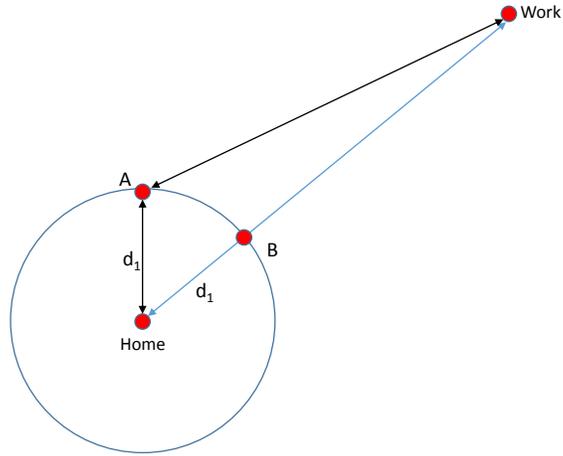


Figure 8: High excess distance region (example 1)

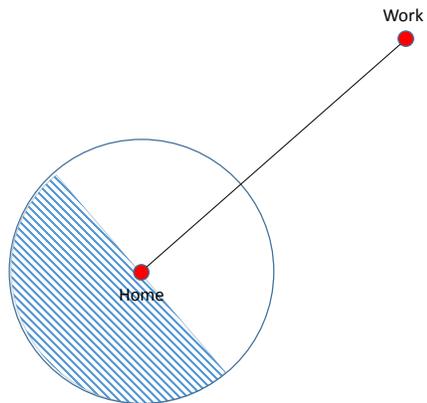


Figure 9: High and low excess distance locations (example 2)

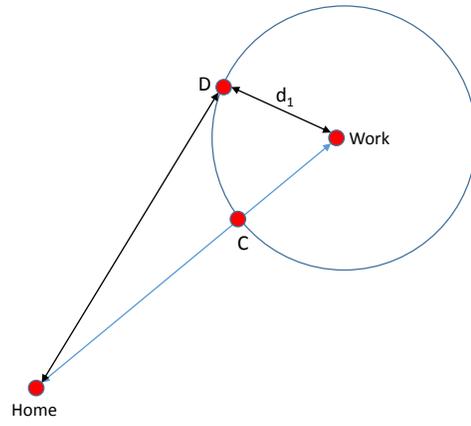


Figure 10: High excess distance region (example 2)

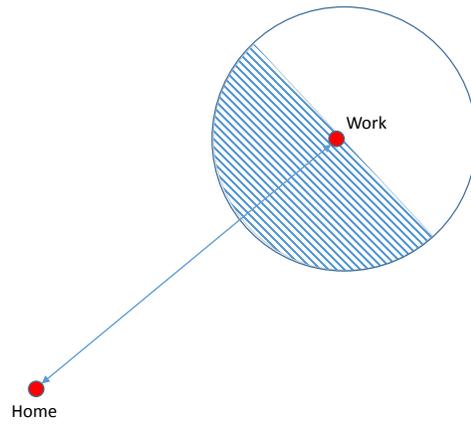
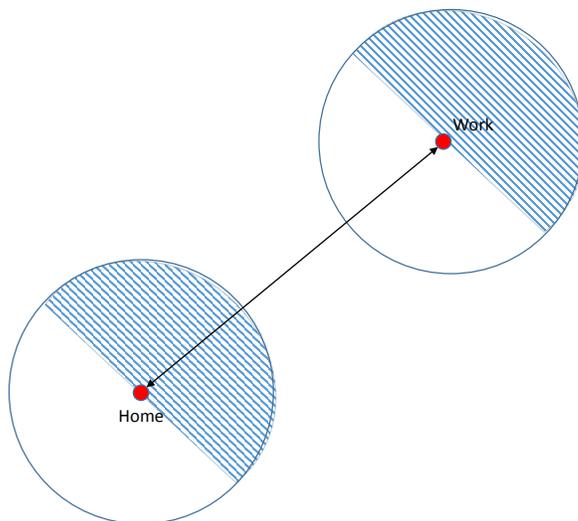


Figure 11: High excess distance regions (combined)



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 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 .

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.

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.³⁷ 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.

Table 6 in the main text summarizes the relationship between my estimated disutility of distance and tract characteristics. I then use the coefficients in this table to adjust the mean estimates of the disutility of distance to account for these observable characteristics at the tract level, using the following equation:

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

where δ_{tc} is my estimate of the disutility of distance for an individual in tract t in city c ; $\bar{\delta}_c$ 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 \bar{S}_c , \bar{D}_c and \bar{N}_c are the city means of these variables.

³⁷For 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.