

# Do segregated neighborhoods cause segregated interactions?

## Evidence from Flickr

Kirsten Cornelson,  
University of Notre Dame\*

### Abstract

Segregated social networks help to perpetuate racial inequality in the presence of peer effects on economic outcomes. We currently have little understanding, however, of the factors influencing segregation in social networks, or how we might reduce social segregation through policy. In this paper, I explore whether reducing residential segregation can reduce social segregation as well. I first show theoretically that desegregation policy can have a *negative* effect on cross-racial interactions, even when sharing a neighborhood positively affects the probability that two individuals interact. The ambiguity arises because desegregation policy may reduce the correlation between neighbors' unobserved characteristics. People have few other-race neighbors, but those neighbors are unusually well-matched to them; this creates inter-racial interactions that may be destroyed if people are induced to live in more diverse, but less well-matched, neighborhoods. Next, I examine what effect desegregation is likely to have in practice. I use my model, along with lower- and upper-bound values of the causal effect of neighborhoods on interactions, to simulate the effect of completely desegregating U.S. cities on interaction rates. I compare the results of my simulation to a measure of the actual inter-racial interaction rate, measured with a new dataset I have created using geocoded Flickr photographs. In all versions of my simulation, inter-racial interactions fall as a result of desegregation. JEL codes: R0,J1,Z1

---

\*3051 Jenkins-Nanovic Hall, Notre Dame, IN, 46616. Phone: 574-631-3461. Email: kcornels@nd.edu

# 1 Introduction

There is a high degree of residential segregation between blacks and whites in the United States. 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 a whole.<sup>1</sup> 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.

A large literature in economics has argued that segregation is harmful for black economic outcomes (e.g., Cutler and Glaeser, 1997; Card and Rothstein, 2007; Ananat, 2011). Sociologists claim that the harmful effect of residential segregation operates, in part, through its effect on social segregation (e.g., Wilson, 1987; Massey and Denton, 1993; Krysan and Crowder, 2017). According to this perspective, residential segregation creates segregated friendship networks through some causal effect of neighborhoods on social interactions. Segregated social networks then help perpetuate racial inequality through channels such as job referrals, social norms, or other kinds of social influence.<sup>2</sup>

While there is a large literature in economics that provides support for the empirical relevance of social influences on economic behavior (e.g., Duflo and Saez, 2003; Bayer, Ross and Topa, 2008; Dahl, Loken and Mogstad, 2014), there is currently little evidence evaluating the first piece of this argument - the impact of residential segregation on social segregation. Would reducing segregation in U.S. cities lead to more inter-racial interactions? If so, how large is this effect likely to be?

In this paper, I attempt to answer these questions by using a transferable utility model of social interactions to examine the theoretical and empirical implications of reducing segregation in U.S. cities. The first contribution of the paper is to show theoretically that policies aimed at reducing segregation can have a *negative* impact on the frequency of inter-racial interactions. This is true even if neighborhoods have a positive causal effect on interaction rates.<sup>3</sup> The reason is that racial segregation arises through a process of residential sorting, which also ensures that neighbors

---

<sup>1</sup>Calculation by author, using 2010 Census data.

<sup>2</sup>An alternative perspective is that segregation is harmful because it restricts access to good schools, job opportunities, or other spatial amenities (e.g. United States Department of Housing and Urban Development (2014)). If this is the case, then policy makers can focus on improving access to these amenities, rather than on reducing segregation itself. In contrast, the social-interactions hypothesis implies that residential segregation must be reduced to equalize opportunities across race.

<sup>3</sup>In my model, I allow everyone in the city to potentially interact with anyone else in the city, and treat neighborhoods as having a causal effect on the utility from each potential choice. This approach combines the effects of neighborhoods on friendship formation (extensive margin) and choosing how much time to spend with each friend (intensive margin.)

are similar to each other along other, non-race dimensions. While a white individual may have relatively few black neighbors, the black neighbors she does have are likely to be unusually good matches for her. Having these particular other-race partners close by causes her to have more black interactions. Policies that reduce segregation are also likely to reduce the correlation in unobserved characteristics between neighbors, because they reduce the incentive for individuals to sort into their otherwise-optimal neighborhoods. If the aforementioned white individual is induced to move to a more diverse neighborhood, she will have more other-race partners nearby, but those potential partners are likely to be less well-matched to her. As a result, it is unclear whether her inter-racial interactions will increase or decrease.

The theoretical analysis highlights that the process of residential sorting has both positive and negative effects on the inter-racial interaction rate. In practice, which of these effects dominates? In the next part of my paper, I answer this question by considering the extreme case of a policy where individuals are randomized to neighborhoods; as a result, there is neither residential segregation, nor any correlation between neighbors' interaction-relevant characteristics. I show that my model can be used to simulate the inter-racial interaction rate that would occur in this world. By comparing the results of my simulation to the actual inter-racial interaction rate, this thought experiment helps us to assess the net contribution of residential sorting to social segregation. While the policy used in my simulation is extreme, it nonetheless helps to quantify the relative magnitudes of the positive and negative forces induced by desegregation policy.<sup>4</sup>

The key input to my simulation exercise is an estimate of the causal effect of location on the probability that two individuals interact. While I do not have an estimate of this parameter, I argue that I can place lower and upper bounds on the range of empirically relevant values for it. The lower bound comes from the simple observation that individuals must incur travel costs in order to meet. Moving two individuals further apart will reduce the value of their interactions because it increases the time the interaction takes.<sup>5</sup> We already have a good sense of how individuals value

---

<sup>4</sup>A key assumption in my analysis is that racial preferences remain fixed after redistributing individuals to neighborhoods. I discuss the potential implications of relaxing this assumption in the discussion of my results at the conclusion of the paper.

<sup>5</sup>Of course, desegregating residential areas will only affect travel time for interactions that start from home, not those that start from work. In the empirical section of the paper, I focus on interactions that take place on weekends, when this is likely to be true. The impact of desegregation policy could be smaller in magnitude if it does not affect workplace segregation; conversely, it could be more positive if workplaces desegregate in response to a change in residential locations. I return to this point in the conclusion.

travel time from a large literature estimating the demand for spatially differentiated goods such as gas stations or coffee shops (e.g., Thomadsen, 2005; Davis, 2006; McManus, 2007). I use estimates from this literature in my model to consider the lower bound case. While I cannot place a precise upper bound on the value of the causal parameter, I show that increasing it beyond about 10 times the lower bound value leads to no discernible changes in predicted behavior. This is because the causal effect becomes sufficiently strong at that point to ensure that individuals interact only with people in their immediate neighborhood; this is true whether the causal effect is 10 times the lower bound value, or 1000 times the lower bound value. Therefore, the results from the lower bound estimate and 10 times this amount cover all empirically relevant cases.

Finally, in order to interpret my simulation results, I need to be able to compare the simulated inter-racial interaction rate to the actual rate of black-white interactions. There are currently no large, publicly available datasets that contain information on inter-racial interaction rates for a broad sample of the population.<sup>6</sup> A final 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 geotagged social media photographs, which I run through face detection and race classification software. I validate this measure using survey data, which show that the racial composition of faces in social media photographs corresponds very closely to the racial breakdown of an individuals' friends.<sup>7</sup> Using this dataset, I document a substantial degree of black-white social segregation in the U.S. The typical non-black Flickr user in my data sees black friends for approximately 3.5% of their interactions. By comparison, this number would be 13.1% in a perfectly integrated world, where interactions reflected the demographics of a user's city.

The results of my simulation suggest that residential sorting has a positive contribution to the frequency of inter-racial interaction. If the cities in my sample were completely desegregated, the fraction of time non-black Americans would spend socializing with black friends would fall from about 0.9% currently<sup>8</sup> to about 0.2-0.3%, depending on the value of the causal parameter I use.

---

<sup>6</sup>Previous research on this topic has relied either on Add Health, a survey of teenagers from the 1990's (Patacchini, Picard and Zenou, 2015), or on specialized proprietary datasets such as email records from Dartmouth University students (Marmaros and Sacerdote, 2006), or early Facebook data (Sara Baker and Puller, 2011).

<sup>7</sup>Throughout the paper, I focus on time spent with friends as my primary measure of interactions; this both is because I can measure this type of interaction behavior using the Flickr data, and because I expect this type of interaction to create the strongest peer effects. Of course, desegregation might have a more positive impact on casual interactions - e.g. chatting with coffee shop employees, or seeing neighbors on the street. I return to this point in the conclusion, when interpreting my results.

<sup>8</sup>This is calculated based on the fact that 25% of white Americans socialize with friends on any given day, and

While the results for inter-racial contact are similar across all simulations, the specific mechanism behind the result does differ somewhat as I change the causal parameter value. If we think that location has a relatively weak effect on the probability that two individuals interact, then the fraction of all interactions that are inter-racial does rise slightly after desegregation - from about 3.5% to 4.1%. This is more than offset, however, by a very steep decline in the overall interaction rate, leading to fewer black-white interactions on net. This result arises because the preferences that rationalize existing interaction behavior in this case point to a very strong preference for ones' current neighbors. Desegregation therefore causes the total number of interactions to fall quite substantially. This effect is stronger for same-race than for different-race interactions, which leads the proportion of inter-racial interactions to rise. Nonetheless, the policy decreases the total amount of inter-racial contact.

If we believe that the causal effect of location on interactions is very high, the story changes somewhat. The parameter estimates from my model in this case indicate that there is currently a relatively weak correlation between neighbors' characteristics. This is particularly true for same-race pairs, implying that same-race partners living in near or far neighborhoods are good substitutes for each other. This is less true for different-race pairs, however, who still show a reasonably strong degree of within-neighborhood complementarity. This pattern means that when individuals are redistributed to different locations in a city, the total interaction rate falls by a more moderate amount than in the case where the causal effect is small. However, of the interactions that are destroyed through redistribution, most of them are different-race pairings. As a result, both the proportion and magnitude of inter-racial interactions fall.

My results suggest that, regardless of what we believe about the causal effect of neighborhoods on interactions, the process of residential sorting tends to lead to more inter-racial interactions on net. If policies designed to reduce segregation in a city also undermine the tendency of neighbors to be similar to each other, they may not achieve their intended goal of reducing social segregation (although other goals, such as improving minorities' access to spatial amenities, may still be served.)

In the next section of the paper, I outline the theoretical model that I use as the basis for my simulation exercise, which is derived from the marriage matching model of Choo and Siow (2006).

---

3.5% of these interactions are with black friends.

I derive the equilibrium conditions, and then explain in the following section how these conditions can be used to simulate the effect of complete desegregation. In the following section, I explain the data I use to get the key inputs into the simulation. I then present the results, and conclude.

## 2 Theory

### 2.1 Model setup

Agents live in neighborhoods in a city, and must decide whether to interact with anyone, and if so, with whom to interact. I assume that each person chooses a single interaction partner. The partner characteristics that the agent chooses are race and location.<sup>9</sup> I refer to agent types as  $(i, r)$ , where  $i$  indexes the individual's location (her neighborhood) and  $r$  indexes their race. I assume that there are a finite number of racial groups  $R$  and neighborhoods  $N$ .

Utility is assumed to be transferable between partners. These transfers act as a price, allowing the market to clear between different groups of potential partners. Let  $\tau_{irsj}$  represent the equilibrium transfer from type  $i, r$  partners to type  $j, s$  partners; note that this quantity may be either positive or negative.

The causal effect of neighborhoods on interactions is assumed to operate through a parameter  $\delta$  that reduces the utility of an interaction when two individuals are moved to neighborhoods that are further apart. I make no assumptions about the channel through which this parameter operates: at a minimum, it should reflect travel costs, but could also include any other influence of location on interaction probabilities. I also make no assumption about whether the causal effect of neighborhood influences who meets who (extensive margin), or how frequently they interact (intensive margin). In my model, the causal effect simply affects the total number of interactions between residents of different neighborhoods.

Putting all of these pieces together, the utility of an individual agent  $g$  of type  $(i, r)$  interacting with an individual  $h$  of type  $(j, s)$  is:

---

<sup>9</sup>In practice, agents may care about a wide variety of partner characteristics, such as age or education. As I explain below, preferences over the race and location of interaction partners can be interpreted as reflecting average values of these other characteristics across different groups, as well as any race- or location- specific preferences.

$$U_{ghirjs} = \alpha_{js}^{ir} - \delta d_{ij} + \tau_{irjs} + \epsilon_{ghirjs}$$

where  $\alpha_{js}^{ir}$  is the average utility generated for agents of type  $(i, r)$  from socializing with an agents of type  $(j, s)$ ;  $d_{ij}$  is the physical distance between the two agents;  $\tau_{irjs}$  is the equilibrium transfer from type  $(i, r)$  partners to type  $(j, r)$  partners; and  $\epsilon_{ghirjs}$  is an I.I.D. shock with a type I extreme value distribution. The term  $\alpha_{js}^{ir}$  will capture average levels of education, age, or other characteristics that affect the match value between a typical member of type  $(i, r)$  and a typical member of type  $(j, s)$ . The term  $\epsilon_{ghirjs}$  can be interpreted as reflecting individual partners' deviations from the average level of characteristics for members of their types; for example,  $\epsilon_{ghirjs}$  may be positive if both partners have higher-than-average education for their tract and race. This term also reflects more idiosyncratic shocks, such as personality characteristics, that affect the valuation of a match.

The agent may also choose not to socialize at all, which I denote as choosing partner type “0”. I normalize the intrinsic utility from spending time alone to zero.

## 2.2 Equilibrium

Following Choo and Siow (2006), the model described above will lead to a quasi-demand function describing the number of  $(j, s)$  type interactions demanded by all individuals of type  $(i, r)$ , which depends on the transfer  $\tau_{irsj}$  (interpreted as the “price” in this model). The quasi-demand function takes the form

$$\ln(\mu_{irsj}^q) = \ln(\mu_{ir0}) + \alpha_{js}^{ir} - \delta d_{ij} + \tau_{irsj} \quad (1)$$

where  $\mu_{irsj}^q$  is the total number of type  $(j, s)$  interactions demanded by agents of type  $(i, r)$  and  $\mu_{ir0}$  is the equilibrium number of  $(i, r)$  agents who choose to spend time alone. In equilibrium, demand for these interactions by agents of type  $(i, r)$  must equal the “supply” of these interactions by agents of type  $(j, s)$ :

$$\ln(\mu_{irsj}^s) = \ln(\mu_{0js}) + \alpha_{ir}^{js} - \delta d_{ij} - \tau_{irsj} \quad (2)$$

The transfers  $\tau_{irsj}$  will adjust to ensure that this is the case. Setting demand for interactions

equal to supply, solving for  $\tau_{irsj}$ , and plugging this back into the quasi-demand equation gives:

$$\ln\left(\frac{\mu_{ijrs}}{\sqrt{\mu_{i0}\mu_{0j}}}\right) = \frac{1}{2}[\alpha_{js}^{ir} + \alpha_{ir}^{js}] - \delta d_{ij}$$

To close the model, note that there is an adding-up constraint. If we denote the population of race  $r$  living at  $i$  as  $f_{ir}$ , then  $\mu_{ir0} + \sum_{t \in R} \sum_{k \in N} \mu_{irkt} = f_{ir}$ . This lets us rewrite the equilibrium equation entirely in terms of the endogenous terms  $\mu_{irjs}$  and the parameters of the model:

$$\ln\left(\frac{\mu_{irjs}}{\sqrt{(f_{ir} - \sum_{t \in R} \sum_{k \in N} \mu_{irkt})(f_{sj} - \sum_{t \in R} \sum_{k \in N} \mu_{ktjs})}}\right) = \frac{1}{2}[\alpha_{js}^{ir} + \alpha_{ir}^{js}] - \delta d_{ij} \quad (3)$$

This equation says that the frequency of interactions between types  $(i, r)$  and  $(j, s)$  (scaled by a function of the total number of partners of each type) is equal to the per-partner surplus created by interactions between these types of individuals.<sup>10</sup> Note that it is only the total surplus that matters here; any difference between  $\alpha_{js}^{ir}$  and  $\alpha_{ir}^{js}$  is irrelevant, since the agent who gets more utility from the interaction can always compensate the other partner through the transfer  $\tau_{irjs}$ . For the purposes of the simulation exercise, I will rewrite the term  $\frac{1}{2}[\alpha_{js}^{ir} + \alpha_{ir}^{js}] = \alpha_{irjs}$ , the average utility generated by the interaction.

This equilibrium equation holds for every  $(i, r), (j, s)$  pair in the city. If we let  $N$  be the number of types, this condition gives us a system of  $\frac{(1+N)N}{2}$  equations. Choo and Siow (2006) show that, given values for the right-hand side of the equilibrium equation and a vector of population supplies  $f_{ir}$  for each neighborhood, there is a unique vector of social interactions  $\mu_{irjs}$  that will solve this system of equations. I use this fact in my simulation exercises, described in the following subsections of the paper.

## 2.3 Implications

To see the implications of the model for desegregation policy, start by rewriting Equation 3 to express the inter-racial interaction frequency of a type  $i, r$  individual:

---

<sup>10</sup>Note that this is identical to the standard result from discrete choice problems that forms the basis for logistic regression models. The only difference is that we have extended the problem to a case where agents must coordinate their decisions. The assumption of transferable utility permits this coordination to take place.

$$\mu_{irs} \equiv \sum_j \mu_{irjs} = \sum_j e^{\alpha_{irsj} - \delta d_{ij}} \sqrt{\mu_{ir0} \mu_{js0}} \quad (4)$$

Next, write the utility generated by an interaction in the following way:

$$\alpha_{irjs} = \alpha_{rs} + \varepsilon_{irs} + \eta_{irjs} \quad (5)$$

The term  $\alpha_{rs}$  represents the average utility created by an interaction between someone of race  $r$  and someone of race  $s$  within the city. The term  $\varepsilon_{irs}$ , which has a mean of zero across all  $i$ , represents the deviation from this average for individuals in neighborhood  $i$ . This term would be negative, for example, if race- $r$  individuals in neighborhood  $i$  were unusually racist against individuals of race  $s$ . The terms  $\eta_{irjs}$ , which also has mean zero across all neighborhood pairs, represents the additional deviation when someone of race  $r$  from neighborhood  $i$  interacts with someone of race  $s$  from neighborhood  $j$ . This will reflect any specific compatibility between these two groups.

Note that both the terms  $\varepsilon_{irs}$  and  $\eta_{irjs}$  arise because of residential sorting, a process that creates systematic differences in observable and unobservable characteristics between individuals living in different neighborhoods. To the extent that these characteristics affect the utility from a match, we expect that the term  $\alpha_{irjs}$  will vary depending on the location of the two interaction partners (i.e. that both  $\varepsilon_{irs}$  and  $\eta_{irjs}$  will typically not be zero.) I assume that both  $\varepsilon_{irs}$  and  $\eta_{irjs}$  are normally distributed, although not in a way that is independent of residential location.

Using this expression in Equation 4 and taking logs gives:

$$\ln(\mu_{irs}) = \alpha_{rs} + \varepsilon_{irs} + \ln(\sum_j e^{\eta_{irjs}} \hat{\mu}_{irjs}) \quad (6)$$

where  $\hat{\mu}_{irjs} = e^{-\delta d_{ij}} \sqrt{\mu_{ir0} \mu_{js0}}$ . Note that  $\hat{\mu}_{irjs}$  can be interpreted as the interaction frequency that would take place if we eliminated preferences over interaction partner characteristics (i.e. set  $\alpha_{rs}$ ,  $\varepsilon_{irs}$ , and  $\eta_{irjs}$  equal to zero for all partner combinations).<sup>11</sup> In other words, it is the number of interactions between types  $(i, r)$  and  $(j, s)$  that would exist if *only* the causal effect of neighborhoods and population supplies influenced the interaction rate.

<sup>11</sup>Note that the terms  $\hat{\mu}_{irjs}$  are *not* equilibrium outcomes. If we eliminated preferences over partner characteristics, the overall interaction rate (and therefore the terms  $\mu_{ir0}$ ) would change, which is not captured in my construction of  $\hat{\mu}_{irjs}$ . We can instead think of  $\hat{\mu}_{irjs}$  as the initial non-equilibrium result of eliminating preferences.

Finally, rewrite the term in brackets as:

$$\begin{aligned}\sum_j e^{\eta_{irjs}} \hat{\mu}_{irjs} &= \frac{1}{N} (\sum_j e^{\eta_{irjs}}) * (\sum_j \hat{\mu}_{irsj}) + \sum_j (e^{\eta_{irsj}} - \overline{e^{\eta_{irsj}}}) (\hat{\mu}_{irsj} - \overline{\hat{\mu}_{irsj}}) \\ &= \beta_{irs} \hat{\mu}_{irs} + \gamma_{irs}\end{aligned}$$

where I have written  $\sum_j \hat{\mu}_{irsj} \equiv \hat{\mu}_{irs}$  in the second line. The inter-racial interaction equation then becomes

$$\ln(\mu_{irs}) = \alpha_{rs} + \varepsilon_{irs} + \ln(\beta_{irs} \hat{\mu}_{irs} + \gamma_{irs}) \quad (7)$$

Write  $\beta_{rs} \equiv E(\beta_{irs})$ . This term should be equal to  $e^{\frac{\theta_{rs}}{2}} > 0$ , where  $\theta_{rs}$  is the variance of  $\eta_{irjs}$ . As a result, we expect  $\beta_{rs}$  to be greater than one. In other words, a 1-unit increase in  $\hat{\mu}_{irs}$  will typically lead to a more-than-one unit increase in actual interactions. This occurs because  $\eta_{irsj}$  is assumed to affect interactions in a non-linear way: a small increase has a bigger impact on interactions than a small decrease.

Similarly, write  $\gamma_{rs} \equiv E(\gamma_{irs})$ . This term will be equal to  $N * Cov(e^{\eta_{irjs}}, \hat{\mu}_{irsj})$ . We should expect this covariance to be positive so long as individuals who live close to one another are more similar, on average. In this case,  $\eta_{irsj}$  will decline in distance: the match value between people who live closer together will typically be higher than the match value between people who live further away. The term  $\hat{\mu}_{irsj}$  also declines in distance, because of the term  $e^{-\delta d_{ij}}$ . As a result, we would expect  $\gamma_{rs}$  to be positive.

The term  $\gamma_{rs}$  captures the positive effect of residential sorting on interactions. The positive covariance between match values and location leads  $\sum_j e^{\eta_{irjs}} \hat{\mu}_{irjs}$  to be larger than  $\hat{\mu}_{irs}$ , the interaction frequency that would prevail in the absence of preferences. In other words, the interaction between the causal effect (captured in  $\hat{\mu}_{irsj}$ ) and sorting (captured in the term  $\eta_{irsj}$ ) create additional inter-racial interactions. Specifically *because* there is a causal effect of neighborhoods on interaction utility, individuals will tend to socialize more when the people they like best also live close by. The process of residential sorting, in which individuals choose neighborhoods on the basis of their unobservable characteristics and preferences, ensures that this is the case.

Using this equation, we can think about the effect of a policy that reduces residential segregation by race - perhaps through a voucher program, or tax subsidies for people living in diverse neighborhoods. The intended effect of the policy is to increase  $\hat{\mu}_{irs}$  for the case where  $r \neq s$  through a reduction in distance between black and white residents of a city. As shown in the equation, this will indeed tend to increase inter-racial interactions if nothing else in the model changes. The policy, however, is not “all else equal”: it is also likely to change the distribution of the  $\eta_{irjs}$ . In particular, if the program is successful, it must be the case that it induces some people to live where they otherwise wouldn’t. For these individuals, either their preferences or their financial circumstances mean that this new neighborhood was not their optimal choice prior to the policy. In contrast, others in the neighborhood *did* choose that location, implying that their preferences or financial constraints are different from those of their new neighbors. The neighborhood will end up being more diverse not only in terms of race, but also along other dimensions as well. As neighborhoods become more diverse, they become less predictive of match values; individuals will be less likely to strongly prefer residents of one neighborhood over another (although their preferences for *individuals* within neighborhoods will remain equally strong.) In my model, this corresponds to a reduction in both  $\beta_{rs}$  and  $\gamma_{rs}$ . Interactions are destroyed by the desegregation policy, because individuals no longer live close to the people they like best. As a result, the net effect of the policy is unclear.

The same logic holds for the desegregation policy I consider in my thought experiment: randomizing individuals to neighborhoods. In contrast to more marginal programs, which will reduce the terms  $\beta_{rs}$  and  $\gamma_{rs}$ , this policy will completely eliminate these terms from the model. My simulation, then, can be thought of as the most extreme version of the general tendency for desegregation to destroy interactions by reducing residential sorting.

### 3 Simulation exercise

In practice, does residential sorting (and the associated racial segregation) tend to increase or decrease inter-racial interactions in U.S. cities? To answer this question, I conduct a thought experiment in which I use the equilibrium conditions of my model to simulate the effect of completely randomizing individuals to neighborhoods. Randomization eliminates both racial segregation, as well as the tendency for neighbors to be similar to each other. By comparing the inter-racial inter-

action rate that occurs in this world to the actual inter-racial interaction rate, I can assess the net effect of residential sorting on social segregation.

The steps of the simulation exercise, which I describe in more detail below, are as follows:

1. Estimate the preference terms  $\alpha_{rs}^c$  from my model.
2. Randomly reassign residential locations.
3. Calculate the frequency of inter-racial interactions that would take place with this new distribution of individuals, using the preference terms estimated in step 1.

### 3.1 Step 1

The first step of my simulation procedure is to estimate the parameters from my inter-racial interactions equation:

$$\ln(\mu_{irs}) = \alpha_{rs} + \varepsilon_{irs} + \ln(\beta_{rs}\hat{\mu}_{irs} + \gamma_{rs})$$

For this estimation, I start with data on the interaction frequency with members of different races at the tract level (described in more detail in the data section). This forms the left-hand side of my regression equation. For each tract  $i$  in my dataset, I then construct the term  $\hat{\mu}_{irs} = \sum_j \hat{\mu}_{irjs}$ , the inter-racial interaction frequency that would be predicted if the causal effect of distance was the only factor influencing interaction decisions. To construct this term, I require several pieces of information: an estimate of the causal effect of neighborhoods on interaction utility ( $\delta$ ), as well as the unmatched population of each neighborhood ( $\mu_{ir0}$ ). Again, I defer my discussion of where I get these pieces of information until the data section.

Once I have constructed the dependent and independent variables, I estimate the terms  $\alpha_{rs}$ ,  $\beta_{rs}$  and  $\gamma_{rs}$  using a non-linear least squares regression across tracts within a city, separately for each race-pair. Because the parameter  $\beta_{rs}^c$  is expected to be equal to  $e^{\frac{\theta}{2}}$ , where  $\theta$  is the variance of  $\eta_{irjs}$ , I restrict this parameter to be greater than or equal to 1.

Figure 1 shows how the parameters are identified in this equation. As shown in panel A), the regression fits a non-linear curve between  $\hat{\mu}_{irj}$  and  $\ln(\mu_{irs})$  in order to minimize the sum of squared

errors.<sup>12</sup> To see how the parameters are identified, panels B) and C) show how the fitted curve changes with  $\alpha$ ,  $\beta$  and  $\gamma$ . As shown in Panel B), increasing the term  $\alpha$  creates an upward shift in the fitted curve, implying an equal percentage increase in interactions at all levels of predicted interaction. The estimate of  $\alpha$  can therefore be interpreted as the average percentage by which actual interactions deviate from predicted interactions in a city.

As shown in Panel C), increasing the term  $\beta$  - moving from the blue line to the green line - will also increase interactions, but does so more at higher levels of predicted interactions. In the context of the model, this is because the variance in  $\eta$  tends to magnify differences in  $\hat{\mu}_{irs}$ . Also shown in Panel C), increasing the term  $\gamma$  - moving from the blue line to the red line - has the opposite effect: it increases interactions but does so more (in percentage terms) at low levels of interactions. In the context of my model, the parameter  $\gamma$  can be thought of as the average number of additional interactions that take place for each tract in a city because of the correlation between  $\eta_{irsj}$  and distance. Because this number is the same for every tract, the *percentage* increase that is generated through this process is larger when  $\hat{\mu}_{irs}$  is small. Together, the terms  $\beta$  and  $\gamma$  therefore jointly govern both the slope and the concavity of the fitted curve. As shown in the top line of Panel C), increasing both terms together has an effect similar to increasing  $\alpha$ . This highlights the fact that it is the relative size of  $\beta$  and  $\gamma$  that is economically interesting, rather than their levels. Cities with high  $\beta$  and low  $\gamma$  will be those where the relationship between predicted and the log of actual interactions is relatively steep and concave; cities with high  $\gamma$  and low  $\beta$  will have a relatively flat and linear relationship between these variables.

## 3.2 Step 2

The goal of this simulation exercise is simulate the frequency of inter-racial interactions that would take place if we eliminated residential segregation completely. The specific counter-factual I use is one where individuals of different races are evenly distributed throughout the city, with each individual being randomly assigned to a specific neighborhood.

To achieve this distribution of individuals, I assume that every neighborhood has the same total

---

<sup>12</sup>Note that these are errors in *percentage* terms, which are different from the level error terms that would be minimized by an OLS regression. An OLS regression would take the form  $\mu_{irs} = e^\alpha \beta \hat{\mu}_{irs} + e^\alpha \gamma + \varepsilon_{irs}$ , and would be unable to separately identify the parameters  $\alpha$ ,  $\beta$ , and  $\gamma$ . In contrast, the regression I estimate can be written as  $\mu_{irs} = e^{\alpha + \varepsilon_{irs}} \beta \hat{\mu}_{irs} + e^{\alpha + \varepsilon_{irs}} \gamma$ . The non-linearity between the error terms and the outcome helps to identify the terms in the model.

population as it does currently, but that the fraction of individuals of race  $r$  is equal to the proportion of these individuals in the *city* population for every tract. Random assignment of specific individuals implies that each neighborhood will, on average, contain a representative sample of individuals of each race.<sup>13</sup> I use this fact in the derivation of my simulation results below.

### 3.3 Step 3

The third step of my exercise is to calculate the inter-racial interaction frequencies that would take place if individuals had the same preferences estimated in the first step of the exercise, but had the geographic distribution simulated in the second step. To do this, I use the equilibrium equation from my model. Letting  $N$  denote the set of neighborhoods in a city, this is:

$$\ln \left( \frac{\mu_{irjs}}{\sqrt{(f_{ir}^* - \sum_{t \in R} \sum_{k \in N} \mu_{irtk})(f_{js}^* - \sum_{t \in R} \sum_{k \in N} \mu_{ktjs})}} \right) = \hat{\alpha}_{rs} - \delta d_{ij}$$

In this equation,  $f_{ir}^*$  is the new number of individuals living in neighborhood  $i$  of race  $r$ , after the redistribution. The preference term on the right hand side is the average match value for an  $r, s$  pair, as estimated in step 1 of the simulation. Importantly, the match values do not vary systematically across neighborhood pairs the way they do in real life. That is, the terms  $\varepsilon_{irs}$  and  $\eta_{irjs}$  will both be equal to zero. This is because we have eliminated any sorting into neighborhoods on the basis of observable or unobservable characteristics. The average value of a match between a person of race  $r$  in neighborhood  $i$  and a person of race  $s$  in neighborhood  $j$  will be the same, no matter which  $i$  and  $j$  we choose.<sup>14</sup> Using my estimates  $\hat{\alpha}_{rs}$ , along with calculations of  $\delta d_{ij}$  and the simulated population supplies  $f_{ir}^*$ , I can solve this system of equations for  $\mu_{irjs}$ . These can then be aggregated into predictions about both the overall interaction rate, and the inter-racial interaction rate.

---

<sup>13</sup>In practice, random assignment will not result in an exactly even distribution of interaction-relevant characteristics across neighborhoods. It is possible for my model to accommodate this random variation. To do this, I can allow my estimated  $\alpha_{rs}$  to vary with the observable characteristics of each individual (e.g. education, age) and the average level of these characteristics among individuals of race  $s$  in the city; this provides an approximation of how interaction values are affected by these characteristics. Then, I could randomly re-assign individuals and construct a predicted  $\alpha_{ijrs}$  for each neighborhood-race pair, based on the characteristics of each neighborhood. This procedure should not affect my results on average, however; as a result, I do not implement it here.

<sup>14</sup>Of course, the match value of specific individuals in these neighborhoods may be higher or lower than this; this is captured by the error terms  $\epsilon_{hgirjs}$  in each individual's utility function.

## 4 Data

To perform my simulation, I need four pieces of information. First, I need to know the existing population of each neighborhood,  $f_{ir}$ , separately by race. Secondly, I need to know the geographic distance between neighborhoods within a city; this corresponds to  $d_{ij}$  in my model. Third, I need an estimate of the causal effect of distance,  $\delta$ . Finally, I need two pieces of information on the social interaction behavior of Americans. First, I need to know how often individuals in different neighborhoods socialize. This, combined with information on population, will allow me to estimate the terms  $\mu_{ir0}$  from my model. Secondly, I need to know how frequently they socialize with members of each race. This will allow me to construct the left-hand side terms  $\mu_{irs}$  in my model. In this section, I describe where I get this information.

### 4.1 Population distribution

Information on the population distribution by neighborhood and geographic distances are available from the U.S. Census Bureau. Throughout the analysis, I will define a neighborhood as a Census tract. The average population in a Census tract in the U.S. is around 4,000 individuals. I restrict the analysis to pairs of Census tracts within the same Core-Based Statistical Area (CBSA).<sup>15</sup> There are 933 CBSAs in the United States (excluding Puerto Rico), of which I use 148 in my main analysis sample.<sup>16</sup> These cities make up just over two-thirds of the U.S. population.

Different measures have been used to capture the degree of racial segregation within cities. One popular measure is the Duncan index, which measures the fraction of black or non-black residents within a city that would have to move to produce an even distribution of racial groups over Census tracts.<sup>17</sup> The first column of Table 1 shows that the mean Duncan index in my sample is 0.527, indicating that about half of all residents in a typical city would have to move to achieve perfect integration. The next rows show how the Duncan index varies across the four Census regions.

---

<sup>15</sup>CBSAs consist of “one or more counties and includes the counties containing the core urban area, as well as any adjacent counties that have a high degree of social and economic integration (as measured by commuting to work) with the urban core.” (Census Bureau, 2016.)

<sup>16</sup>The restrictions I impose in order to be able to estimate my non-linear regressions are that I must have an estimate of the disutility of travel, which requires information on both travel speeds and hourly wages, and that I must have at Flickr users represented in at least 25 tracts. These restrictions reduce the pool of CBSAs to 148, with most of the eliminations coming from the second restriction.

<sup>17</sup>The Duncan is calculated using the formula  $D_c = \sum_t \left| \frac{\text{Black}_{tc}}{\text{Black}_c} - \frac{\text{Nonblack}_{tc}}{\text{Nonblack}_c} \right|$  where  $\text{Black}_{tc}$  is the number of black individuals living in tract  $t$  in city  $c$ , and  $\text{Black}_c$  is the total number of black individuals in the city (and similarly for non-black individuals).

According to this measure, segregation is highest in the Midwest, with an average Duncan index of 0.616, and lowest in the Pacific, with an average Duncan index of 0.458.

## 4.2 Distance

To measure the geographic distance between pairs of Census tracts, I use shapefiles provided by the U.S. Census Bureau. I calculate the great-circle distance between the central latitude and longitude of each pair of tracts within a CBSA. The third column of Table 1 shows that, on average, two randomly selected tracts within the same CBSA are 44.1 km (27.4 miles) apart. This varies from 38.9 km (24.2 miles) in the Pacific to 46.4 km (28.8 miles) in the Northeast.

## 4.3 The causal effect of distance

In my model, the parameter  $\delta$  captures the causal effect of distance on interaction utilities. While I do not know the value of  $\delta$ , this parameter should at a minimum be equal to the disutility of travel. I therefore use estimates of the disutility of travel as lower bound estimates of  $\delta$  in my simulation exercise. I then explore the impact of increasing  $\delta$ .

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.<sup>18</sup>

Table 2 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

---

<sup>18</sup>For example, Manuszak and Moul (2009) use a tax hike near Cook County to estimate consumer's willingness to pay to travel across county lines to purchase gasoline.

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). Because I am interested in setting a lower bound on the parameter  $\delta$  from my model, I start by assuming that individuals value their time at the hourly wage; the case where the valuation of time is higher is covered by my alternative estimates.

To construct an estimate of the disutility of travel for each city in my sample, I start by estimating the hourly wage for each city. Information on hourly wages is available for some metropolitan areas from the Bureau of Labor Statistics. In order to preserve the majority of CBSAs in my sample, however, I instead impute average hourly wages by using information on state-level hourly wages and the ratio of median income in the CBSA to median income in the state. For each city, I construct a “per minute” disutility of travel equal to either one or two times the wage per minute in that city.

Next, I convert the dollar valuation of time to utility terms using the estimates from Houde (2012), which are available in both units of measure. This gives me an estimate of disutility per minute, which I then convert into “per kilometer” format by using information on travel speeds from the Google Maps API.<sup>19</sup>

I have sufficient information on income and travel speeds to calculate the disutility of travel for 862 CBSAs, which include all 148 CBSAs in my main analysis sample. As shown in the fourth column of Table 1, the mean disutility of travel across cities is around 0.469 per kilometer, which corresponds to a dollar valuation of around \$0.47 per kilometer in 2010 dollars. The table also shows how the disutility of travel varies by region. The cost of traveling 1 km is highest in the Northeast and Pacific, and lowest in the South.

A limitation of this procedure is that variation across cities is imposed by assumption, not by revealed behavior. We can get some sense of whether the implied distaste for travel actually corresponds to individuals’ travel behavior 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

---

<sup>19</sup>Specifically, I choose 10 randomly selected pairs of Census blocks within a CBSA and query the API for a driving time between them on a Saturday afternoon at 3 pm. In my main results, I use Flickr photos taken on weekends to measure social interactions, in order to capture time periods when individuals are likely to be leaving from home.

how individuals move throughout their home cities. In the Appendix, I show that my estimated disutility of travel significantly predicts the fraction of photographs that are taken within 5 and 10 km of a Flickr user’s estimated home location; the coefficients are positive, but not significant, for photos taken within 1 and 3 km of home.<sup>20</sup> This suggests that my estimates are picking up actual travel behavior, although they capture long-distance travel better than shorter-distance travel.

## 4.4 Social interactions

My simulation exercise requires that I have two pieces of information about social interactions: the overall interaction rate, and the frequency of inter-racial interactions, by race and by neighborhood. Unfortunately, this information is not available in any large, publicly available dataset. The 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 has information on either residential location or cross-racial interactions; however, both pieces of information are available only for a relatively small subsample of respondents.<sup>21</sup> Additionally, this data source only allows us to measure interaction behavior for a very specific group (teenagers), nearly two decades ago.

To measure interaction behavior, I instead rely on a combination of the American Time Use Survey (ATUS) and a novel dataset I have constructed using Flickr photographs. In the remainder of this section, I provide more information on how I construct measures of interaction behavior from these sources of data.

### 4.4.1 Interaction rates

I use the American Time Use Survey (ATUS) to arrive at my estimates of the overall interaction rate for each tract/race pair. The ATUS is an annual survey of a representative sample of Americans

---

<sup>20</sup>If I expand the sample to include all 862 CBSAs for which I have estimates of the disutility of travel and any Flickr users, the coefficients are positive and highly significant at all distances.

<sup>21</sup>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.

that asks respondents to keep a diary recording what they are doing and who they are with for every 15 minute segment of the day. My measure of interaction rates in the ATUS will be the probability that a respondent spends any time with friends on his or her diary day. This measure corresponds closely to the decision margin in my theoretical model. The mean of this variable is 22.7% for both races, indicating that about one-fifth of the population spends time with friends on a randomly selected day. This varies somewhat across different times of the week, with an average of 20.6% on weekdays and 24.5% on weekends. Because I will be focusing on the impact of desegregating residential locations, which are more likely to impact social interactions on weekends, I will focus on weekend socializing in both the ATUS and Flickr results.

Unfortunately, the ATUS does not contain geographic information below the state level. To arrive at ATUS estimates for tracts, I use the demographic information present in the ATUS to predict the interaction rate based on tract demographics. Details of this procedure are included in the Appendix. My predicted interaction rate varies from 18.5%-50.1% for blacks (with a mean of 25.3%), and from 13.7% to 65.2% for whites (with a mean of 24.5%). Because the interaction rate is based on demographic characteristics available in the American Community Survey, I am able to predict it for nearly all of the approximately 70,000 tracts in the U.S. Table 1 shows how the interaction rate varies by region.

To validate the use of my predicted interaction rates, I compare the ATUS predictions of interaction rates to actual interaction rates in a survey of MTurk workers. Details of the survey are provided in the Appendix. The key survey measures for the purposes of validating the interaction rate from ATUS are a series of time use questions that refer to activities in the previous day. I ask respondents what they were doing for each 3 hour chunk of the day, and who they were with. Using this series of questions, I construct an indicator for whether the respondent spent any time with friends on the day in question. 25.8% of the workers in my survey spent time with friends the previous day, which is quite similar to the ATUS mean.<sup>22</sup>

MTurk surveys automatically include information on the respondents' latitude and longitude while taking the survey. Assuming that this location will typically be the user's home, I use these locations to connect respondents to a predicted tract-level interaction rate constructed from the

---

<sup>22</sup>The survey was run on a Sunday-Tuesday, with the majority of respondents answering on a Monday. Note that the questions referred to the previous day, which means I am capturing interaction behavior primarily on Sundays.

ATUS data. Then, I examine whether the ATUS prediction corresponds well with their actual interaction behavior.<sup>23</sup>

A regression of the MTurk worker’s interaction indicator on his or her predicted interaction rate produces a coefficient of 0.787, which is significant at the 1% level. This coefficient is not statistically distinguishable from 1, which is the coefficient I would expect in this regression. The constant is 0.046 and is not significantly different from 0. When combined with the fact that there is a great deal of error in my “home” assignment processes (which should bias the coefficient in this regression downward), this regression suggests that the estimated interaction rates do a good job of predicting actual interaction behavior.

#### 4.4.2 Inter-racial interactions

To measure the inter-racial interaction rate, I rely on a new dataset I have created 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). I downloaded a large sample of public photographs on Flickr, along with their metadata, and ran them through face detection and race classification algorithms. The racial breakdown of faces in the photographs provides with information on individuals’ inter-racial interaction rates. The metadata, which contain information such as the make of the camera and the time of the photograph, sometimes include “geotags”, which are latitude and longitude coordinates appended by cameras that have access to the internet (smart phones, for example, and higher-end digital cameras.) These geotags allow me to link Flickr users to cities and neighborhoods. I assign each user to the modal CBSA in which they take pictures, and to their modal Census tract within the CBSA. More details about both the race classification algorithm and the home assignment are available in the Appendix.

Several concerns may arise with the use of Flickr photographs to measure inter-racial interactions and/or residential location, including the accuracy of the racial classifier; whether Flickr users are a representative sample; whether users’ photographs are an accurate reflection of actual social behavior; and whether home locations can be inferred from photograph locations. In the Appendix, I provide detailed evidence designed to assuage these concerns. Because measurement of inter-racial

---

<sup>23</sup>Unfortunately, I have too few black respondents (113) to be able to examine the relationship between these variables in the black sample; I therefore restrict myself to the non-black sample in this exercise.

interactions is a central contribution of the paper, I also summarize this evidence here:

*How accurate is the racial classifier?*

I assess the accuracy of the face detection and race classification algorithms by hand coding a set of 17,500 photographs and comparing my results to the results of the algorithms. The face detection algorithm has a high rate of false negatives (missing about 50% of faces - primarily those that are very small), and an 8% rate of false positives (mostly animals, statues, and art). The detection algorithm does not appear to be racially biased. The race classification algorithm has higher accuracy, at 85% for whites and 75% for blacks. Because the two groups are not equally represented in the data, this error nonetheless creates bias in my point estimates of the fraction of black faces. I correct for this by fitting a non-linear model between the fraction of black faces found by the algorithm and the actual fraction of black faces in the hand coded data. Using this adjustment on the entire set of photographs, I find a relative black interaction rate of around 3.7%.

*Are Flickr users a representative sample?*

Using the demographics of the assigned home location, I show that Flickr users differ on several dimensions from the U.S. population: they are located in bigger cities and in wealthier and more educated tracts. I attempt to correct for these observable differences by regressing the fraction of black faces on tract demographics, and adjusting my estimates to match the demographics of the U.S. population. This results in a slightly lower estimated black interaction rate - 3.5%, vs. 3.7% in my Flickr sample. Table 1 shows that this fraction is very similar by region. To the extent that we believe the Flickr sample differs from the population on unobservable characteristics, however, my estimates may not be valid outside of the sample.

Additionally, based on my inspection of the hand-coded photographs, it appears that very few users are black: less than 1% of the users in the hand-coded sample have a majority of black faces in their photographs. Because of this, I treat my entire Flickr sample as non-black. As I discuss further in the next section, this does not interfere with my simulation procedure. Every cross-racial interaction for a non-black person represents a cross-racial interaction for a black person as well. I assume that these interactions are generated by the preferences of both interaction partners, and attempt to estimate a parameter capturing this joint surplus.

*Are users' photographs an accurate reflection of their actual social behavior?*

It is possible that the rate of black faces in a Flickr user’s photographs is not representative of the user’s actual rate of interaction with black friends. This could arise if users “curate” their photographs, for example. To examine the relationship between social media photographs and actual social behavior, I turn to my survey of MTurk workers. In this survey, I ask detailed questions about respondents’ social contacts and behavior; I also ask for information about the racial breakdown of faces in a recent social media photograph. I then compare the user’s actual social behavior to the behavior depicted in the photo. As shown in the Appendix, the relationship between the reported fraction of interactions that occur with black friends and the fraction of black faces in the users last social media photograph is almost perfect: a regression indicates that the correlation between the two is close to 1. For a population that uses social media, the fraction of black faces in photographs is an excellent measure of actual social behavior.

*Can we infer someone’s home location from the location of their photographs?*

In the Appendix, I provide three pieces of evidence that I am correctly identifying the home tract of Flickr users, on average. First, the home tract is visited far more often than other tracts: a typical user is observed in her home Census tract on 14 separate days over the course of a year, but only 3 times in any other tract that she visits. Secondly, I link my photographs to the Foursquare database to show that there are fewer venues (bars, restaurants, offices, etc) of any kind in the immediate vicinity of the user when she is in her home tract. Taken together with point 1, this suggests that users spend more time in their home tracts despite the fact that there is less reason for them to be there. Finally, I show that there is a stronger link between the ethnic breakdown of a user’s home tract and her own ethnicity (as measured by last names) than for other tracts in her home city.

## 5 Results

Step 1 of my exercise is estimating the residual parameters  $\alpha_{ww}$  and  $\alpha_{wb}$  from my model. Recall from Section 3 that the estimating equation I use for this is:

$$\ln(\mu_{irs}) = \alpha_{rs}^c + \ln(\beta_{rs}^c \hat{\mu}_{irs} + \gamma_{rs}^c) + \varepsilon_{hirs} \quad (8)$$

In this regression,  $i$  indexes Census tracts,  $r$  and  $s$  are racial categories, and  $c$  is a city. The

dependent variable is the log of the estimated number of daily interactions of each type that occur at the tract level. This is calculated as the tract population multiplied by interaction rate for non-blacks (measured in ATUS), multiplied by the fraction of interactions for non-blacks that occur with either black or non-black partners (measured in the Flickr data). The key variable on the right-hand side of the equation is  $\hat{\mu}_{irs}$ , the predicted interaction rate that would occur if only the causal effect of distance affected interaction rates. I construct this term based on the disutility of travel, distance, and population supplies for every tract and link these estimates to Flickr users based on their home tract locations. I run this regression separately using non-linear least squares for white-white interactions and white-black interactions, and separately for each city.

Because the term  $\beta_{rs}^c$  is equal  $e^{\frac{\theta^2}{2}}$  in expectation, where  $\theta$  is the variance of  $\eta_{irsj}$ , I expect this term to be greater than 1, and constrain this to be the case. This procedure produces a separate estimate of  $\alpha_{ww}^c$  and  $\alpha_{wb}^c$  for each city  $c$ , and, when combined with the error term  $\varepsilon_{hirs}$ , for each tract that contains a Flickr user. In a next step, I adjust the estimates of  $\alpha_{ww}^c$  and  $\alpha_{wb}^c$  to take account of the observable differences between Flickr users and other residents of their cities. Specifically, I run a regression of  $\varepsilon_{hirs}$  on a set of tract observable characteristics, combining the data from all cities within a region, and use the coefficients to predict the average level of  $\varepsilon_{hirs}$  for non-black individuals in the city; I then add this to the term  $\hat{\alpha}_{sr}^c$  to arrive at my final estimates.<sup>24</sup>

Note that because the terms  $\hat{\mu}_{irs}$  will differ depending on the assumed value of  $\delta$ , I will arrive at different values of  $\alpha_{ww}^c$  and  $\alpha_{wb}^c$  for each value of  $\delta$  I use. The results of my simulation procedure therefore differ both because the causal effect of reallocating people to new locations changes with  $\delta$ , and because the preferences that rationalize their current behavior will be different for different levels of  $\delta$ .

Table 3 shows the key inputs and the average results of the non-linear regressions, for different levels levels of  $\delta$ . The data on the left-hand side of the equation - the actual interaction frequency for each tract - is the same regardless of which value of  $\delta$  I use. Based on my interaction data, I estimate that non-black people in a typical tract have around 1,300 interactions with other non-black people each day.<sup>25</sup> The interaction frequency that I would predict in the absence of any

<sup>24</sup>The specific characteristics I use are: tract population, density, median age, median income, proportion of the population that is white, black, Hispanic, and Asian, the proportion of the population that has less than high school, high school, some college, and a bachelor's degree, and city fixed effects. Tables are available upon request.

<sup>25</sup>The average Flickr user lives in a tract that has a population of about 5,000 non-black people. Approximately 25% of these interaction on any given day, with the majority of interactions occurring with other non-black people;

interaction preferences is much higher, ranging from 25,273 when  $\delta$  is very low to 3,536 when  $\delta$  is at its highest level.<sup>26</sup> This is the interaction rate that would occur if all interactions produced the same intrinsic utility as staying at home alone.<sup>27</sup> Because the actual interaction frequency is lower than the predicted level, the estimates of  $\alpha_{ww}$  are negative. In other words, because people choose to interact less frequently than they choose to spend time not interacting, it must be the case that interacting generates less utility than not interacting, on average.<sup>28</sup>

The number of predicted interactions falls as I assume higher and higher levels of  $\delta$ . This is because distance is assumed to destroy the utility from interactions, at a rate that goes up as  $\delta$  increases. As the predicted interaction rate falls, people’s behavior looks closer to the predicted level; as a result, the estimates of  $\alpha_{ww}$  decline in magnitude. Note that there is very little change in either the predicted interaction level or the estimates of  $\alpha_{ww}$  when I increase  $\delta$  beyond 100 times the disutility of travel. Even if the causal effect of distance is higher than this, it will make very little difference empirically.

The second panel of Table 3 shows the same information for  $\alpha_{wb}$ . Unsurprisingly, the estimated preference for cross-racial interactions is substantially more negative than the preference for same-race interactions. As shown in the last row of the table, the ratio of  $\alpha_{wb}$  to  $\alpha_{ww}$  is around 2 for the case where  $\delta$  is equal to the disutility of travel. This indicates that inter-racial interactions reduce utility by about twice as much as same-race interactions, compared to staying home. Less obviously, the estimated preference gap between same- and different-race interactions grows as  $\delta$  increases. This is because the terms  $\alpha_{ww}^c$  fall at a faster pace as  $\delta$  increases, relative to  $\alpha_{wb}^c$ . In other words, when it comes to same-race interactions, we can justify the interaction rate either by assuming a general distaste for interactions and a relatively low cost of distance, or by assuming that people are relatively indifferent between interacting and not interacting in general ( $\alpha$  close to zero), but that the costs of interacting with people far away are very high. When it comes to different-race interactions, however, there is no cost of distance that can entirely justify the low level of interaction.

Table 3 also shows the average level of  $\beta_{ww}^c$ ,  $\beta_{wb}^c$ ,  $\gamma_{ww}^c$ , and  $\gamma_{wb}^c$  that are estimated for each this results in about 1,300 interactions per tract.

<sup>26</sup>Recall that  $\hat{\mu}_{irs}$  is not an equilibrium object; as a result, it is not constrained to be less than or equal to the population of the tract. When  $\delta$  is low, the number of predicted interactions far exceeds the tract population.

<sup>27</sup>Recall that this utility is normalized to zero in my model.

<sup>28</sup>Note, however, that this is not true *for interactions that actually take place*. Recall from my model that individuals receive a “shock” to the value of interacting on any given day. When this shock is sufficiently positive, individuals will choose to interact in spite of their negative baseline utility.

level of  $\delta$ . When  $\delta$  is very small, the relationship between predicted and actual interactions is quite weak, for both same- and different-race interactions:  $\beta$  is small, and  $\gamma$  is large. As  $\delta$  increases, the estimated relationship between predicted and actual interactions becomes much stronger. Figure 2 and Figure 3 show these patterns visually, plotting both the data and the curves that are estimated if I run the non-linear regression on all cities simultaneously. When  $\delta$  is low, the relationship between  $\hat{\mu}_{irs}$  and  $\ln(\mu_{irs})$  is close to linear for both same- and different-race interactions. At the highest level of  $\delta$ , the relationships are steeper and more concave.

Recall that  $\gamma_{rs}^c$  increases interactions at the origin, with a relatively smaller effect at higher levels of predicted interactions. Effectively, a high estimate of  $\gamma_{rs}^c$  - which is what we get for both same- and different-race interactions in the case where we assume  $\delta$  is low - says that individuals with low predicted interactions have many more interactions than we would otherwise expect. In the context of my model, this is generated by the fact that we like our neighbors: even people with relatively few interaction partners around them will choose high interaction rates, because the people nearby are the ones they want to partner up with. If we redistribute people, we undo this correlation between match value and geographic location, and eliminate the additional interactions generated by  $\gamma$ . This will tend to push the overall interaction rate down. The effect of redistribution on inter-racial interactions depends both on the causal effect of moving black and white individuals closer together, as well as on the relative sizes of the terms  $\gamma_{ww}^c$  and  $\gamma_{wb}^c$ .

Step 2 of my exercise is to randomly reassign individuals to locations in a city, so that each Census tract is the same on both observable and unobservable characteristics. I use this fact when I calculate the equilibrium interaction rates in Step 3, using my equilibrium condition (repeated here for convenience):

$$\ln \left( \frac{\mu_{irjs}}{\sqrt{(f_{ir}^* - \sum_{t \in R} \sum_{k \in N} \mu_{irkt})(f_{sj}^* - \sum_{t \in R} \sum_{k \in N} \mu_{ktjs})}} \right) = \hat{\alpha}_{rs} - \delta d(i, j)$$

Using my estimates  $\hat{\alpha}_{rs}^c$  from Step 1, and the population vectors  $f_{ir}^*$  from Step 2, I can use this equation to calculate the equilibrium frequency of interactions that occurs between any neighborhood-race pair. I then aggregate these predictions into three objects: i) a total interaction rate for each race (defined as the number of interactions involving members of that race, divided by the population); ii) the percentage of interactions for members of each race that take place with members of

the opposite race; and iii) the inter-racial interaction rate.

Table 4 shows the results of this simulation exercise for the non-black population. In the case where  $\delta$  is at its lower bound, redistributing individuals across the city has a very strong and negative impact on the total interaction rate: it is predicted to fall from 25.1% to just 7.1%. The main reason for this is because we have eliminated interactions produced by the parameters  $\gamma_{ww}^c$  and  $\gamma_{wb}^c$ . In Step 1, I estimated that many interactions took place because individuals had a strong preference for the particular people who lived close by. When we randomize individuals to neighborhoods, we eliminate these additional interactions. As a result, the overall interaction rate falls quite substantially.

When  $\delta$  is low, the intervention is successful in raising the *relative* proportion of inter-racial interactions. The percentage of all non-black interactions that involve black people nearly doubles from 3.5% to 6.3%. This occurs for two reasons. First, there is a positive causal effect of moving black and non-black people closer together. Secondly, the effect of sorting seems to be stronger for same-race interactions, relative to different-race interactions ( $\gamma_{ww}^c > \gamma_{wb}^c$ , on average.) Redistribution eliminates many same-race interactions, but a smaller number of inter-racial interactions. Despite the positive effect on the proportion of inter-racial interactions, the total number of inter-racial interactions goes down. The probability that a non-black person interacts with a black person on a given day falls from around 0.9% to around 0.3%. This suggests that the effect of eliminating residential sorting is stronger than the causal effect.

As shown in the next four columns of the table, the impact of desegregation on the interaction rate is less severe when  $\delta$  is higher. This is both because individuals must have a higher intrinsic preference for socializing (higher  $\alpha_{ww}^c$  and  $\alpha_{wb}^c$ ) in these cases, and because the parameters  $\gamma_{ww}^c$  and  $\gamma_{wb}^c$  are smaller. Note however, that the drop is still fairly substantial: at the highest level of  $\delta$ , the non-black interaction rate is around 21.4%, a nearly 20% decline from its actual level of 25.1%. Regardless of what we believe about the magnitude of  $\delta$ , the residential sorting process appears to be a significant input into the general social interaction rate in a city.

While the total interaction rate declines by less in the case of high  $\delta$ , the fraction of interactions that take place with black partners *falls* in these cases. Depending on the level of  $\delta$ , it ranges from 1.0-2.2%, compared to 3.5% currently. When combined with the decline in the total interaction rate, this leads the inter-racial interaction rate to fall quite substantially. This occurs because there

is an asymmetry in the parameters  $\gamma_{ww}^c$  and  $\gamma_{wb}^c$ , which changes as  $\delta$  increases. At low  $\delta$ ,  $\gamma_{ww}^c$  is estimated to be higher for same-race interactions than for opposite race interactions. As a result, eliminating this term reduces same-race interactions more than it reduces inter-racial interactions; this, combined with the causal effect of redistributing individuals, causes the inter-racial interaction rate to rise. At high  $\delta$ , the parameter  $\gamma_{ww}^c$  falls to nearly zero, while the parameter  $\gamma_{wb}^c$  remains positive and large. Eliminating these terms through desegregation causes inter-racial interactions to fall, while there is no similar effect for same-race interactions. While this is offset by the causal effect of moving black and non-black people closer together, the sorting effect dominates, causing the inter-racial interaction rate to fall.<sup>29</sup>

Table 5-Table 6 show the same set of results, broken down separately by region. Table 5 shows that the regions differ substantially in predicted same-race interactions when  $\delta$  is low, but not when  $\delta$  is high. The difference when  $\delta$  is low is driven by differences in city size and density across regions. The similarity at high  $\delta$  occurs because at the highest values of  $\delta$ , individuals are predicted to interact almost entirely within their own tract. Census tracts have similar populations by construction, which means that the predicted interaction rate will not vary much geographically. As a result of these patterns, the terms  $\alpha_{ww}^c$ , which capture the difference between predicted and actual behavior, vary substantially across regions when  $\delta$  is low, and vary little across regions when  $\delta$  is high. This is also true of the terms  $\alpha_{wb}^c$ , but to a lesser extent. There remains some variance in predicted black interactions across regions, because of differences in the fraction of the population that is black. The relative distaste for black interactions (captured by the ratio of  $\alpha_{wb}^c$  to  $\alpha_{ww}^c$ ) is highest in the South and Midwest, with the difference increasing slightly at high levels of  $\delta$ .

Table 6 shows the results of the simulation exercise separately by region. The overall decline in the interaction rate after desegregation when  $\delta$  is low is particularly strong in the Northeast. This is because, as shown in Table 5, the estimated values of  $\gamma_{ww}^c$  are quite high in this region. Interpreting this result through the lens of my model, this suggests that the residential sorting process leads to a particularly strong correlation between match values and location in this region. The positive effect of desegregation on the relative black-white interaction rate is concentrated in the Midwest, and to a lesser extent, the Northeast.

---

<sup>29</sup>Similar results for the black population can be derived if we assume that  $\alpha_{ww}^c = \alpha_{bb}^c$ ; in other words, that black people have similar same- and different-race preferences as non-black people. Because these are symmetric to the non-black results, they tell a very similar story, which is why I do not present them here.

Other than these points, the pattern of results looks strikingly similar across regions. At low  $\delta$ , the other three regions experience large declines in interaction rates and modest declines in the proportion of interactions that are inter-racial. When  $\delta$  is high, the reduction in interaction rates is more muted, but the proportion that are inter-racial falls substantially. The net effect on the total number of inter-racial interactions is similar across all regions and values of  $\delta$ , ranging from -0.6 to -0.8 percentage points, on a base of approximately 0.9%.

## 6 Discussion and Conclusion

This paper highlights an important, and previously unrecognized fact: that residential sorting can *increase* inter-racial contact, even when it results in racial segregation. Despite the fact that neighborhoods must have some causal effect on interactions (if only through the channel of travel costs), my estimates suggest that integrating cities would have a *negative* effect on the number of cross-racial interactions. This occurs because the existing process of residential sorting encourages cross-racial interactions by ensuring that individuals pay a low “price” for interacting with the other-race individuals they like best.

Of course, my results ignore several channels of adjustment that may occur in the long run. First, I have assumed that preferences for same- and different-race interactions remain fixed when individuals are induced to move to different neighborhoods. Would relaxing this assumption change my results? At first glance, the answer would appear to be no: if desegregation lowers inter-racial contact, then, if anything, it is likely to reinforce existing preferences for same-race contact as time goes on. However, it is important to note that my results only capture intentional interactions between adult friends. If more casual interactions that occur within the context of a neighborhood are more strongly affected by desegregation, and affect people’s preferences, then the long-run effect of desegregation could be more positive than the short-run effect I have estimated. Similarly, if desegregation has a bigger impact for children - who may be placed in more diverse school environments as a result - than it may help reduce same-race preferences over time. Therefore, while my results suggest that desegregation would have a negative effect in the short-term, this could be either strengthened or attenuated by preference changes in the long run.

My model also does not consider any effect of residential desegregation on workplace segregation.

I would expect any causal effect of neighborhoods to act primarily on interactions that occur outside of the workplace context. If desegregation leaves these interactions unchanged, then the qualitative nature of my results will be similar: desegregation will reduce inter-racial interactions, although the magnitude will be smaller. However, if people change their work locations in response to desegregation, then workplaces may also become more integrated. This would suggest that the long-run response may be more positive than the short-run response I have estimated.

While desegregation policy may serve many worthwhile goals, including more evenly distributing access to good schools and other public services, my model and empirical exercise show that they may also have unintended consequences for the pattern of social interactions in cities - at least initially. As my analysis highlights, these consequences may not be unambiguously positive, and could actually serve to reduce interactions between members of different social groups. If social interactions have consequences for economic behavior (and a large body of evidence suggests that they do), then desegregation could end up being harmful for minority outcomes.

## 7 Author affiliations

Kirsten Cornelson is an Assistant Professor of Economics at the University of Notre Dame.

## References

- Ananat, Elizabeth Oltmans.** 2011. “The Wrong Side(s) of the Tracks: The Causal Effects of Racial Segregation on Urban Poverty and Inequality.” *American Economic Journal: Applied Economics*, 3(2): 34–66.
- 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.
- Berinsky, Adam J., Gregory A. Huber, and Gabriel S. Lenz.** 2012. “Evaluating Online Labor Markets for Experimental Research: Amazon.com’s Mechanical Turk.” *Political Analysis*.

- Brueckner, Jan K., and Ann G. Largey.** 2008. "Social Interaction and Urban Sprawl." *Journal of Urban Economics*, 64(1): 18–34.
- Card, David, and Jesse Rothstein.** 2007. "Racial Segregation and the Black-White Test Score Gap." *Journal of Public Economics*, 91(11-12): 2158–84.
- Choo, Eugene, and Aloysius Siow.** 2006. "Who Marries Whom and Why." *Journal of Political Economy*, 114(1): 175–201.
- Cutler, David M., and Edward L. Glaeser.** 1997. "Are Ghettos Good or Bad?" *Quarterly Journal of Economics*, 112(3): 827–872.
- Dahl, Gordon B., Katrine V. Loken, and Magne Mogstad.** 2014. "Peer Effects in Program Participation." *American Economic Review*, 104(7): 2049–2074.
- 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.
- 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).
- Houde, Jean-Francois.** 2012. "Spatial Differentiation and Vertical Mergers in Retail Markets for Gasoline." *American Economic Review*, 102(5): 2147–2182.
- Huff, Connor, and Dustin Tingley.** 2015. "Who are these people?' Evaluating the demographic characteristics and political preferences of MTurk survey respondents." *Research and Politics*, 2(3): 1–12.

- 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.
- Krysan, Maria, and Kyle Crowder.** 2017. *Cycle of segregation: social processes and residential stratification*. New York:Russell Sage Foundation.
- 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.
- 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.”
- Sara Baker, Adalbert Mayer, and Steven L. Puller.** 2011. “Do more diverse environments increase the diversity of subsequent interaction? Evidence from random dorm assignment.” *Economics Letters*, 110: 110–112.
- 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.** 2014. "Expanding Opportunity Through Fair Housing Choice."

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

## 8 Figures and Tables

### 8.1 Tables

Table 1: Summary statistics by region

	% black, population	Duncan index	Avg. dist. b/w tracts	Disutility of travel (per km)	Interaction rate	Black interaction rate (% of all int., non-black)
All	11.7%	0.527	44.1 km	0.469	25.7%	3.4%
Northeast	7.7%	0.573	46.4 km	0.533	25.3%	3.6%
Midwest	10.7%	0.616	40.7 km	0.423	26.0%	3.4%
South	19.1%	0.509	42.4 km	0.430	24.9%	3.5%
Pacific	3.7%	0.458	38.9 km	0.511	25.9%	3.3%

This table shows descriptive statistics for the 84 cities in my main analysis sample, and separately by region. The black population and the Duncan index are computed directly from Census data. The average distance between tracts was calculated using shapefiles provided by the U.S. Census Bureau. For details on the estimation of the other columns, please see the Data section.

Table 2: 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 <sup>&amp;</sup>	1
McManus (2007)	Coffee shops, University of Virginia	2000	0.10	0.5-1 <sup>#</sup>
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 <sup>@</sup>	0.75-2.50 <sup>@</sup>
Seim and Waldfogel (2013)	Liquor stores, Pennsylvania	2005	0.46	1.95

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

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

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

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

Table 3: Estimation of preference parameters: regression inputs and results, larger sample of cities

	Assumed value of $\delta$ (relative to disutility of travel):				
	1	2	10	100	1000
Actual number of non-black interaction (tract-level)	1,277	1,277	1,277	1,277	1,277
$\hat{\mu}_{ww}$	25,273	9,361	3,657	3,536	3,536
$\alpha_{ww}$	-4.83	-2.71	-1.45	-1.37	-1.37
$\beta_{ww}$	1.07	1.15	1.43	1.38	1.38
$\gamma_{ww}$	$3.14e^7$	$1.16e^6$	43	60	60
Actual number of black interactions (tract-level)	43	43	43	43	43
$\hat{\mu}_{wb}$	7,543	2,460	778	745	745
$\alpha_{wb}$	-9.29	-8.15	-6.70	-6.57	-6.56
$\beta_{wb}$	1.05	1.08	1.10	1.13	1.12
$\gamma_{wb}$	$1.36e^8$	$5.44e^7$	$1.64e^7$	$1.60e^7$	$1.60e^7$
Relative distaste for black interactions	1.92	3.01	4.62	4.80	4.79

This table shows the average level of the key variables in my non-linear regression model, and the resulting estimated  $\alpha$ 's, for different assumed levels of  $\delta$ .

Table 4: Simulation results, non-black population

		Assumed value of $\delta$ (relative to disutility of travel):				
		1	2	10	100	1000
<b>Total interaction rate</b>						
	Actual	25.1%	25.1%	25.1%	25.1%	25.1%
	Simulated	7.1%	13.1%	19.9%	21.3%	21.4%
	Difference	-18.0pp	-12.0pp	-5.2pp	-3.8pp	-3.7pp
<b>% of interactions black</b>						
	Actual	3.5%	3.5%	3.5%	3.5%	3.5%
	Simulated	6.3%	2.2%	1.1%	1.0%	1.0%
	Difference	2.8pp	-1.3pp	-2.4pp	-2.5pp	-2.5pp
<b>Black interaction rate</b>						
	Actual	0.9%	0.9%	0.9%	0.9%	0.9%
	Simulated	0.4%	0.3%	0.2%	0.2%	0.2%
	Difference	-0.5pp	-0.6pp	-0.7pp	-0.7pp	-0.7pp

This table shows the results of my simulation exercise for the non-black population, for different assumed levels of  $\delta$ .

Table 5: Regression inputs and results, by region

	Northeast		Midwest		South		Pacific	
	1	1000	1	1000	1	1000	1	1000
Value of $\delta^*$								
Actual number of non-black interaction (tract-level)	1,187	1,187	1,289	1,289	1,308	1,308	1,306	1,306
$\hat{\mu}_{ww}$	45,310	3,419	16,898	3,432	15,701	3,535	29,548	3,709
$\alpha_{ww}$	-5.73	-1.41	-6.53	-1.23	-4.10	-1.36	-4.02	-1.46
$\beta_{ww}$	1.08	1.30	1.06	1.20	1.07	1.38	1.06	1.54
$\gamma_{ww}$	$2.57e^7$	128	$1.22e^8$	83	$1.05e^7$	56	98,062	6
Actual number of black interactions (tract-level)	42	42	43	43	44	44	43	43
$\hat{\mu}_{wb}$	16,534	657	4,778	558	5,423	1,051	5,905	576
$\alpha_{wb}$	-10.22	-6.67	-11.40	-7.32	-8.72	-6.76	-8.00	-5.69
$\beta_{wb}$	1.08	1.09	1.03	1.04	1.07	1.25	1.03	1.03
$\gamma_{wb}$	$1.40e^8$	$1.68e^7$	$4.52e^8$	$2.17e^7$	$4.43e^7$	$2.15e^7$	$3.67e^7$	$4.09e^6$
Relative distaste for black interactions	1.78	4.73	1.74	5.95	2.12	4.97	1.99	3.90

This table shows the average level of the key variables in my non-linear regression model, and the resulting estimated  $\alpha$ 's, for different regions and for the lowest and highest assumed levels of  $\delta$ .

\* relative to the disutility of travel

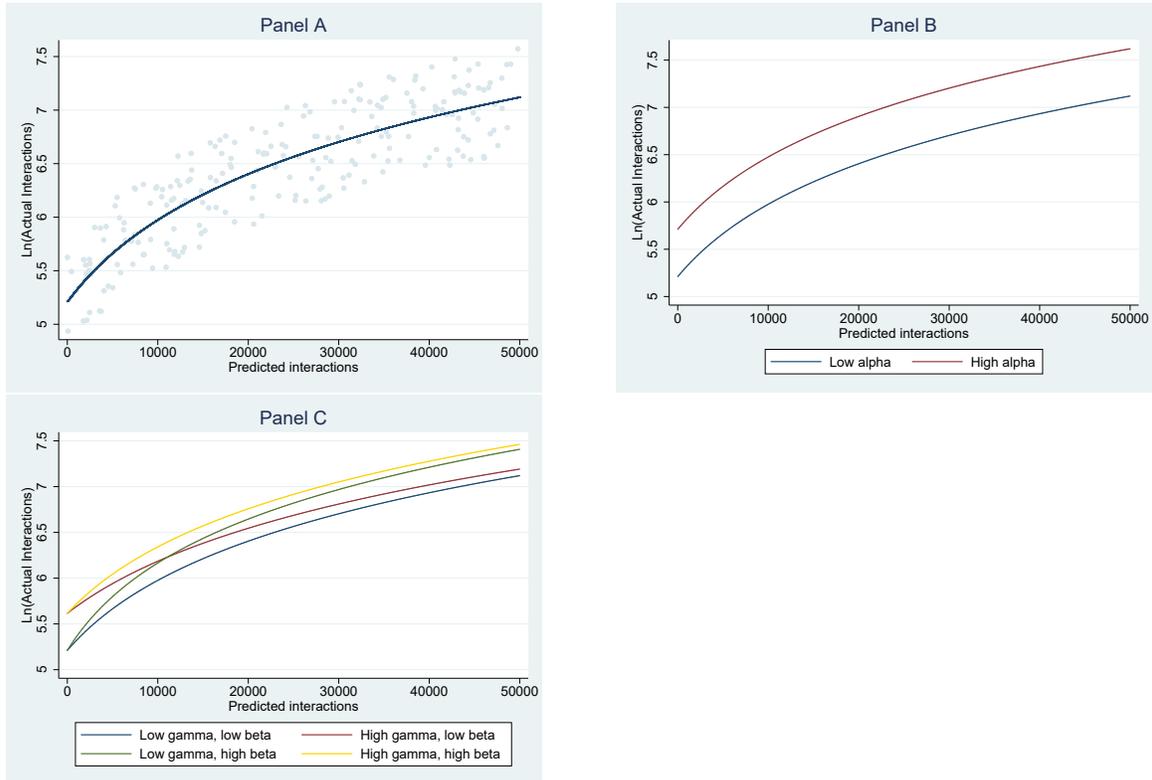
Table 6: Simulation results by region, non-black population

	Northeast		Midwest		South		Pacific	
Value of $\delta^*$	1	1000	1	1000	1	1000	1	1000
<b>Total interaction rate</b>								
Actual	24.2%	24.2%	25.5%	25.5%	25.0%	25.0%	25.7%	25.7%
Simulated	2.8%	22.1%	8.7%	22.9%	9.6%	21.4%	6.6%	20.0%
Difference	-21.4pp	-2.1pp	16.8pp	-2.6pp	-15.4pp	-3.6pp	-19.1pp	-5.7pp
<b>% of interactions black</b>								
Actual	3.6%	3.6%	3.3%	3.3%	3.5%	3.5%	3.4%	3.4%
Simulated	7.5%	1.3%	16.2%	0.7%	3.4%	1.2%	2.0%	0.7%
Difference	3.9pp	-2.3pp	12.9pp	-2.5pp	-0.1pp	-2.3pp	-1.5pp	-2.8pp
<b>Black interaction rate</b>								
Actual	0.9%	0.9%	0.8%	0.8%	0.9%	0.9%	0.9%	0.9%
Simulated	0.2%	0.3%	1.4%	0.2%	0.3%	0.3%	0.1%	0.1%
Difference	-0.7pp	-0.6pp	0.6pp	-0.6pp	-0.6pp	-0.6pp	-0.8pp	-0.8pp

This table shows the results of my simulation exercise for the non-black population, by region, for the lowest and highest values of  $\delta$ .

## 8.2 Figures

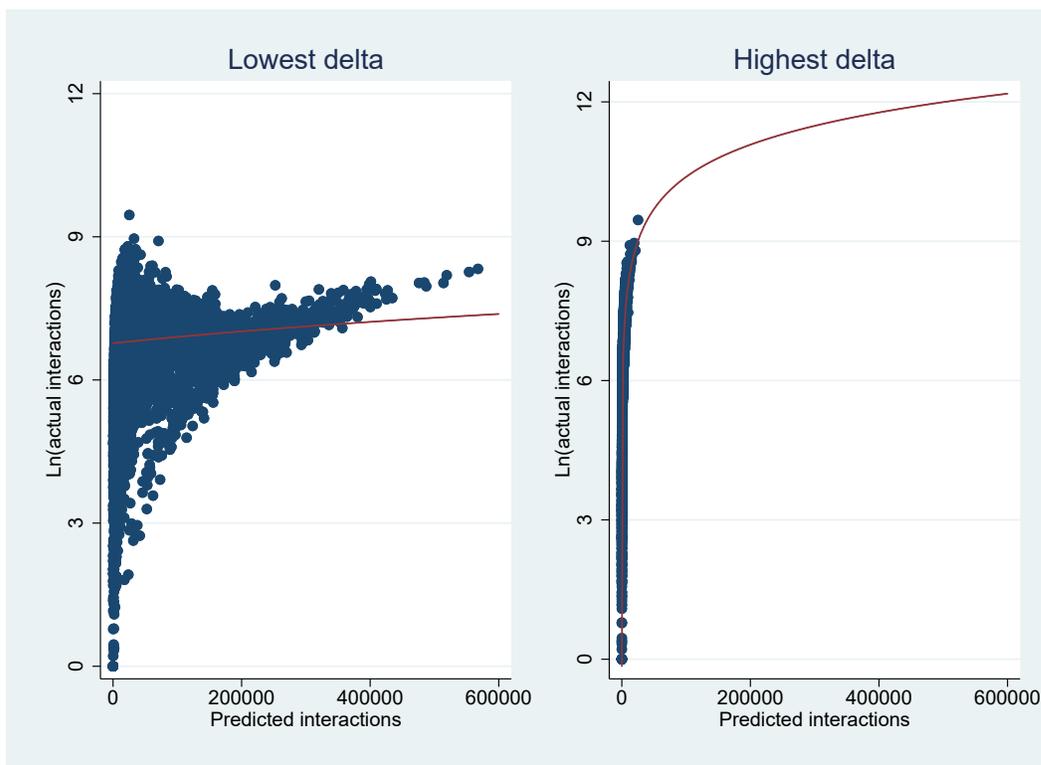
Figure 1: Identification of model parameters: example



These figures show how the parameters  $\alpha_{rs}^c$ ,  $\beta_{rs}^c$ , and  $\gamma_{rs}^c$  are identified. Panel A shows how the choice of parameters fits a curve that minimizes the sum of squared errors. Panel B shows how the estimated curve moves if we increase  $\alpha_{rs}^c$ . Panel C shows the result of increasing  $\beta_{rs}^c$  and  $\gamma_{rs}^c$ , one at a time or together.

lat

Figure 2: Identification of model parameters: white-white interactions



This figure shows the curves that are estimated to best fit the same-race interaction data when my non-linear regression is run for all cities at once. When  $\delta$  is assumed to be low (equal to the disutility of travel), the estimated regression parameters are  $\alpha_{ww} = -6.70234$ ,  $\beta_{ww} = 1.00019$ , and  $\gamma_{ww} = 709,244$ . The resulting curve is shown on the lefthand side of the figure. When  $\delta$  is assumed to be high (equal to 1000 times the disutility of travel), the estimated regression parameters are  $\alpha_{ww} = -1.22980$ ,  $\beta_{ww} = 1.10792$ , and  $\gamma_{ww} = 3$ . The resulting curve is shown on the righthand side of the figure, using the same scale as the lefthand side figure to facilitate comparison of the curves.

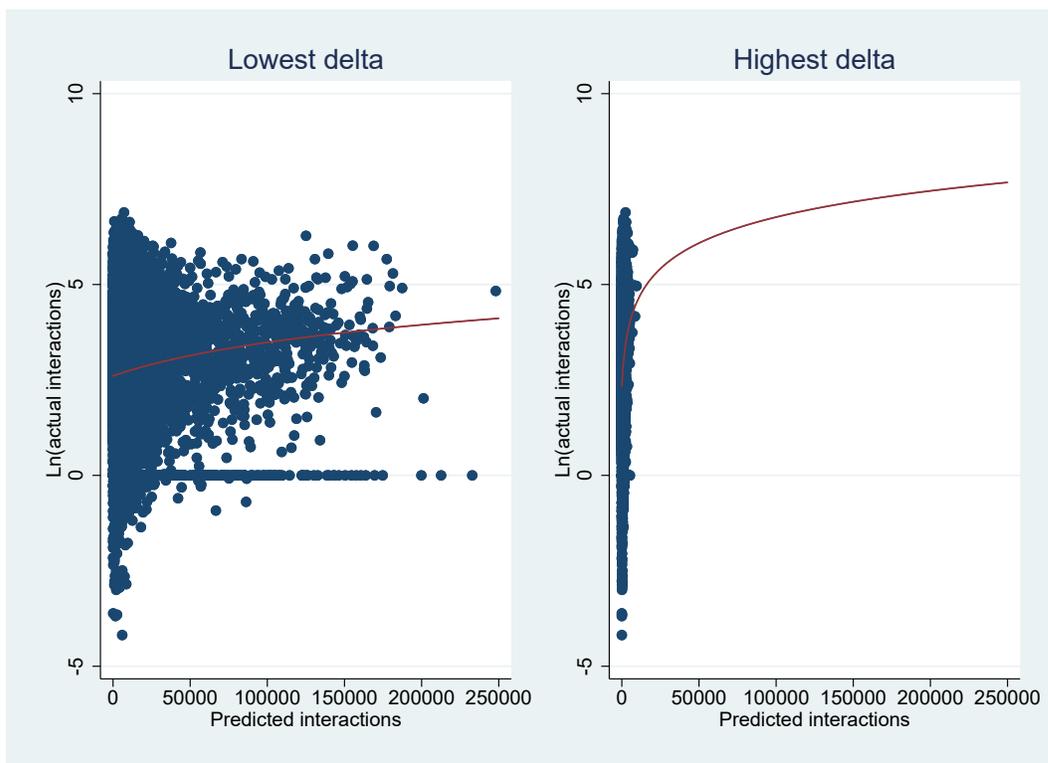
## 9 Appendix

### 9.1 Relating the disutility of travel to Flickr user's travel behavior

Construction of the Flickr dataset is described in a later subsection. For now, note only that the photographs I use are geotagged, which allows me to observe Flickr users in a variety of locations. To test whether my estimates of the disutility of travel at the city level correspond to actual behavior, I examine the relationship between my predicted disutility of travel at the city level and the fraction of photographs that are taken within 1, 3, 5 and 10 km of a Flickr user's estimated home location.<sup>30</sup>

<sup>30</sup>I describe how I infer users' home locations, and provide evidence that I am correctly identifying these locations, in a later section.

Figure 3: Identification of model parameters: white-black interactions



This figure shows the curves that are estimated to best fit the different-race interaction data when my non-linear regression is run for all cities at once. When  $\delta$  is assumed to be low (equal to the disutility of travel), the estimated regression parameters are  $\alpha_{ww} = -8.58900$ ,  $\beta_{ww} = 1.02325$ , and  $\gamma_{ww} = 72,773$ . The resulting curve is shown on the lefthand side of the figure. When  $\delta$  is assumed to be high (equal to 1000 times the disutility of travel), the estimated regression parameters are  $\alpha_{ww} = -4.78580$ ,  $\beta_{ww} = 1.02714$ , and  $\gamma_{ww} = 1,234$ . The resulting curve is shown on the righthand side of the figure, using the same scale as the lefthand side figure to facilitate comparison of the curves.

Table 7 shows the result of this exercise. The table shows that cities with a higher estimated distaste for travel have a higher proportion of photos taken close to home, but that this is only statistically significant for longer distances. The coefficients on  $\delta$  for photos taken within 1 or 3 km of home are positive but non-significant. The coefficient on  $\delta$  for photos taken within 5 km of home is 0.128 and significant at the 10% level, and the coefficient on photos taken within 10 km of home is 0.099 and is significant at the 5% level. To interpret these latter coefficient, note that a one-standard deviation in  $\delta$  (approximately 0.105) is associated with a 1.2 percentage point increase in the fraction of photos

taken within 5 km of home, which is a 1.4% increase over the mean of 88.5%.<sup>31</sup>

## 9.2 Inferring interaction rates from tract demographics

I first run a regression of the probability of socializing in ATUS on various demographic characteristics. Table 8 shows how the probability of weekend socializing varies with age, education, and geographic region, separately for black and non-black respondents. For both racial groups, the interaction rate shows a U-shape in age; the coefficients on age and age squared indicate that interactions decline with age until approximately age 58-61, before starting to rise again. For the non-black sample, the interaction rate shows an approximately linear and increasing pattern in education. For the black sample, interaction rates are similar for all educational groups except for those with post-graduate degrees, who have higher interaction rates. There are no regional differences in interaction behavior for blacks, and only one marginally significant different for whites (with slightly higher interaction rates in the Midwest, compared to other regions of the country.)

I use the results from the model shown in Table 8 to predict the interaction rate for each Census tract based on tract demographics, separately for the black and non-black samples.

I validate the interaction rate produced in this way using a survey of MTurk workers. Details on this survey and the validation are provided in a following subsection.

## 9.3 Validation of ATUS and Flickr measures using an MTurk survey

The survey was administered to approximately 1,600 MTurk workers in the summer of 2018. To be included in the survey, a worker had to be living in the United States, and must have posted at least 1 photograph to social media over the past year. This latter restriction was imposed because I will use the survey to validate the use of my Flickr measure of inter-racial interactions. The survey contained modules asking the workers about basic demographic information; their time use over the previous day, including information on who was present at each moment and the race of those individuals; and their social media use. The demographic module was always presented first, while

---

<sup>31</sup>In principle, disutility of travel may vary at the tract level. This may occur because individuals value their time differently, or because travel speeds differ across tracts. In an earlier version of this paper (Cornelson, 2017), I attempt to estimate tract-level disutilities of travel using information on travel patterns from Flickr. While these parameters did a better job of predicting Flickr users' travel behavior, they made essentially no difference to the interaction results. For this reason, I maintain the assumption of a city-level disutility of travel here.

the other three modules were presented in random order.

Table 9 shows demographic information on the MTurk sample, compared to the U.S. adult population. As might be expected, MTurk workers are not representative of the U.S. population at large: they are significantly younger, more highly educated, and more likely to be white or Asian than other U.S. adults. This is in line with previous work that examines the characteristics of MTurk workers (Berinsky, Huber and Lenz, 2012; Huff and Tingley, 2015). To the extent that these characteristics are reflected in the MTurk workers' residential locations, however, the predicted interaction rates should still be valid for this sample.

## **9.4 Measuring inter-racial interactions in Flickr**

To build this dataset, I began by identifying a set of around 170 million geotagged Flickr photographs, all taken within the U.S. between 2006-2015. To do this, I started by pulling a random sample of about 10% of all geotagged photographs taken in the U.S. over this period. Then, I pulled every photograph ever taken by the approximately 365,000 users in this initial sample. In order to remain in the sample, a Flickr user had to i) take the majority of their photographs in the U.S., ii) post photos taken on at least 3 separate days within a single year, for at least one year, and iii) have at least one face in the sample of photos that I use. The second restriction is required in order to infer a home location for each user. The third restriction is required in order to infer something about the user's inter-racial interaction behavior.

### **9.4.1 Inferring home locations**

I link users to home locations by assigning them to the modal CBSA in which they take pictures, and to their modal Census tract within the CBSA. I do this separately for every year in which a Flickr user posts photographs. In order to abstract from potential moves by Flickr users (some of which may be induced by error in the home assignment process), I assign Flickr users consistently to the home location from the year on which the user posted on the maximal number of days. I also keep only photographs that are taken in a user's home city. I additionally restrict analysis to the CBSAs that contain at least 1 Flickr user in at least 25 separate tracts. This restriction is required in order to run the regression in Step 1 of Simulation B, which identifies the preference parameters

$\alpha$  using cross-tract variation. My final sample of Flickr users is comprised of around 84,000 users.

Appendix Tables 10, 11 and 12 provide evidence that I have correctly identified users' home locations. Table 10 shows that the home tracts are visited far more often than any other tract. The table shows the number of unique "visits" (day by tract level observations) to the home location and to other Census tracts the user visits. The average user is observed in her assigned home Census tract on about 14 separate days; for any other Census tract that the user visits at least once, the mean number of visits is around 3. For a typical Flickr user, about 40% of all visits are in the home tract.

In Table 11, I show that the surroundings in the home tract are observably different from other tracts the user visits. The table shows the types of venues that appear in the home location and in other visited tracts, using information from the Foursquare database. Foursquare is a service that allows individuals to "check-in" at different locations, providing information to friends and family about where they are. Foursquare maintains a database of venues, which is searchable by latitude and longitude. For a sample of around 22,000 owners, I randomly select one photograph taken in their home tract and one photograph taken outside of their home tract, and search the Foursquare database for venues within a 25 meter radius around the location where the photograph was taken.<sup>32</sup> I divide venues into five categories: food and drink (e.g., restaurants, bars, coffee shops), entertainment (e.g., parks, movie theaters, art galleries), stores, offices, and "other". The latter category is mainly comprised of other commercial buildings that are not designated specifically as office buildings; among the most common types of venues in this category are banks, doctor's offices, and barbers/salons. 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. There are fewer venues of all type near the user when she is in her home tract, and the difference is highly significant for four of the five venue types. When combined with the results for visits, these results show that Flickr users spend substantially more time in their "home" tracts, even though there are fewer commercial venues to visit in these areas.

Finally, in Table 12, I use the one piece of information I have on Flickr users - their names - to examine how city and tract demographics correlate with the user's demographics. For users with

---

<sup>32</sup>The Foursquare API maintains rate limits which limit the number of searches to their database each day. This is why I use a smaller sample of owners and photographs.

last names on their profiles, I use information on the racial distribution of the 1000 most common last names in 2010 (Census Bureau, 2010) to construct a probability that the user is white, black, Hispanic or Asian. I then examine how this probability predicts the proportion of people that are of the same race in a user’s CBSA; in the other tracts she visits (aside from the home tract); and in her home tract. Table 12 shows the results for each race in separate panels. The table shows that home tract demographics are more strongly correlated with a user’s race than the demographics of other tracts she visits, or than the demographics of the city as a whole. Increasing the probability that a user is white by 1 percentage point increases the proportion white in her CBSA by 0.073 percentage points; the proportion white in visited tracts by 0.100 percentage points; and the proportion white in her assigned home tract by 0.108 percentage points. The results are similar for other races. This provides further evidence that user’s assigned home tracts are likely to be strongly correlated with their actual neighborhood of residence.

#### **9.4.2 Flickr user demographics**

The information in Table 12 suggests that we can use the characteristics of a Flickr user’s assigned home tract as a proxy for their own characteristics. This provides a way to examine the demographic characteristics of the Flickr sample and how these differ from those of the typical American. Table 13 shows the average characteristics of Flickr users’ CBSAs/tracts, and compares these to those of an average population living in one of the CBSAs in my sample, and to the population of the entire U.S. The Flickr users are more concentrated in larger cities, with an average city size of around 5.5 million, compared to 5.0 million for a typical resident of the same cities. This implies that Flickr users are disproportionately concentrated within the larger cities in my sample. My sample cities are much larger than those inhabited by an average American, which have an average size of 3.8 million. City-level segregation is quite similar for Flickr users and other Americans. However, the two groups live in different areas within cities: Flickr users are concentrated in wealthier and more educated tracts, compared to both other residents of the same cities and the U.S. population in general. Flickr users have a lower proportion of Blacks and Hispanics in their tracts than the typical American but a higher proportion of Asians. While Flickr users are not representative of the U.S. population, I will attempt to use information from the Flickr sample to predict interaction rates

outside of the sample. This process is described in detail below, after I explain my measures of social interactions.

### 9.4.3 Measuring interactions behavior

Once I have linked users to home locations, I measure their inter-racial 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. However, the rate of false negatives appears to be substantially higher in the Flickr data. In a sample of around 17,500 photographs which I hand-coded, I found 13,560 faces in total, while the algorithm found just 6,861. The lower accuracy rate may be due to the fact that the Flickr photographs are often of relatively low quality, and include many faces that are partially turned away from the camera. The rate of false positives also appears to be elevated relative to that reported in Kazemi and Sullivan (2014), at around 8%; this is due primarily to the fact that Flickr users often take pictures of statues, art with faces, and animal faces. Other non-face objects are rarely identified as faces. 18.6% of the photos in my database are found by the algorithm to have any faces in them. Based on the error rates cited above, I estimate that the actual frequency is about twice as large, in the 35-40% range.

I next run all photographs with faces in them through a race classification algorithm, which classified each face as either black or non-black. The algorithm itself was provided as part of the Scikit Learn machine learning module for Python. I trained the classifier using the faces in a random sample of 20,000 Flickr photographs. Table 14 shows the “confusion matrix” from the testing process, which shows the fraction of non-black/black faces that are categorized as non-black or black. Non-black faces are correctly categorized 87% of the time, while black faces are correctly categorized 75% of the time. These error rates are in line with the standards in the literature for this type of classification task (Han and Jain, 2014).<sup>33</sup> While the classifier is correct in the large majority of cases, the error in the assignment process does create a problem for estimating the fraction of black faces in the sample. This is because most of the faces in Flickr photographs are non-black: in a sample of hand-coded photographs, I estimated the fraction of black faces to be around 5%. As a

---

<sup>33</sup>Accuracy rates are much higher in “constrained” classification tasks, where pose and illumination are constant across subjects.

result, the 13% of non-black faces that are classified as black are numerically a far larger set than the 25% of black faces that are misclassified. This results in a much higher estimated frequency of black faces - around 21% - than is actually the case.

To correct for this error, I fit a model linking the proportion of black faces found by the classifier to an actual frequency of black faces according to a set of hand-coded photographs. I first divided Flickr users into 100 groups based on the fraction of black faces found by the classifier (0-1%, 1-2%, etc.) I sampled up to 5 users from each percentile range (not all ranges contained 5 users), then downloaded a random sample of 50 photographs with faces for each of these users. I hand-coded the photographs to arrive at an actual frequency of black faces for each user.

The rate of black faces in this sample is very low. The mean number of black faces is around 5%, and just 14 users have a majority of black faces in their photographs. This is in spite of the fact that users with a high proportion of black faces found by the algorithm were oversampled when I selected photographs to hand-code. It appears then that black individuals are highly under-represented among Flickr users. For this reason, I will assume going forward that my Flickr sample is entirely non-black. As I discuss further in the next section, this does not interfere with my simulation procedure. Every cross-racial interaction for a non-black person represents a cross-racial interaction for a black person as well. I assume that these interactions are generated by the preferences of both interaction partners, and attempt to estimate a parameter capturing this joint surplus.

Figure 4 shows the plot of actual black faces in the hand-coded sample against the number of black faces found by the classifier, with users grouped into 2.5-percent ranges (0-2.5%, 2.5-5%, etc) based on the algorithm's classification. There is an upward sloping and non-linear relationship between the proportion of black faces found by the detector and the user's actual proportion of black faces. Even for users with a very high proportion of black faces found by the algorithm, the actual fraction of black faces is relatively low, with a maximum of just over 30% for users who are found to have 95% of black faces by the algorithm. This supports my assumption that the vast majority of users in my sample are non-black. The line on the graph shows the fitted relationship, which I estimate using a linear and quadratic term. I use this relationship to predict the proportion of black faces for each Flickr user, based on the algorithm's results. The mean fraction of black faces in the broader sample is around 3.7%.

Figure 5 shows how the fraction of black faces found in a user's photographs is related to the percentage of black people in the user's assigned home tract. The relationship between these two variables (shown by the red line) is positive and significant, indicating that individuals living in tracts with more black people have a higher frequency of black interaction. Note, however, that the relationship is quantitatively quite small. For individuals assigned to black-majority tracts, black faces make up an average of 4.2% of all faces in photographs, compared to 3.4% for individuals assigned to black minority tracts. Individuals in all tracts are predicted to have a proportion of black faces under 6%. As indicated by the size of the circles in the graph, the number of individuals living in highly black tracts is quite small; the vast majority of Flickr users are assigned to tracts with fewer than 15% black residents. This raises the possibility that measurement error in either the home assignment or face detection process could be biasing the relationship between tract characteristics for highly black tracts. As these tracts make up a small portion of my overall sample, however, I do not expect this error to affect my estimation procedure.

It is possible that the rate of black faces in a Flickr user's photographs is not representative of the user's actual rate of interaction with black friends. This could arise for two reasons. First, it is likely that users take pictures of family, which means that the racial representation of people in the photographs will tend to overstate the degree of social segregation among friends. Second, it is possible that users "curate" their photographs to show either a higher or lower frequency of inter-racial interactions. To examine the relationship between social media photographs and actual social behavior, I turn to my survey of MTurk workers. In this survey, I ask detailed questions about respondents' social contacts and behavior; I also ask for information about the racial breakdown of faces in a recent social media photograph. I then compare the user's actual social behavior to the behavior depicted in the photo.

Recall from Table 9 that the sample of MTurk workers is disproportionately young and educated, and more likely to be white or Asian than the U.S. population at large. For education and race, this approximately mimics the demographics for the Flickr sample that I estimated in Table 13. The relationship between social photos and social interactions in the MTurk sample is therefore likely to be similar to that found in the Flickr sample.

For each MTurk respondent, I construct a measure of the inter-racial interaction rate based on

the time use portion of the survey. Recall that in this module, I ask users to account for their activities for each 3 hour portion of the day. I also ask who was present during each activity, and for the racial breakdown of the people involved. I use this information to construct the average proportion of black people present when the user spends time with friends alone.<sup>34</sup> This measure corresponds closely to the inter-racial interaction rate in my model, because it captures the relative amount of time spent with black friends, compared to the total amount of time spent with friends. A drawback of this measure, however, is that it can be constructed for a small number of people in my sample: only 461 of the approximately 1500 respondents spent any time with friends on the sample day, and only 217 of these spent time with friends alone (which is required for me to use the fraction of black people present as a measure of the proportion of black friends.)

To connect inter-racial interactions to social media photos, I ask respondents to choose the last photograph of a social event that they posted to a social media site, and to report the racial breakdown of people in the photograph. This exercise is randomly determined to occur before or after the reporting of social behavior. For non-black respondents, the mean reported fraction of black faces is 4%, which is very similar to my Flickr data.

Figure 6 shows a scatter plot of the inter-racial interaction rate for the non-black sample against the fraction of black faces in their social media photographs. The figure shows the fitted line from a regression of the friends measure on the photos measure (with no constant), along with the 45-degree line. The two lines are indistinguishable from one another. The regression coefficient is 1.002 and is significant at the 1% level. The fraction of black faces in this *single* photograph can explain approximately 39% of the variation in the proportion of black friends. This suggests that the fraction of black individuals in a user's photographs is an excellent measure of the user's inter-racial interaction rate.

## 9.5 Correcting for differences between the Flickr sample and the U.S. population

My simulation strategy requires that I know the actual black interaction rate for non-black individuals in my 148 CBSAs. To correct for observable differences between my Flickr sample and other

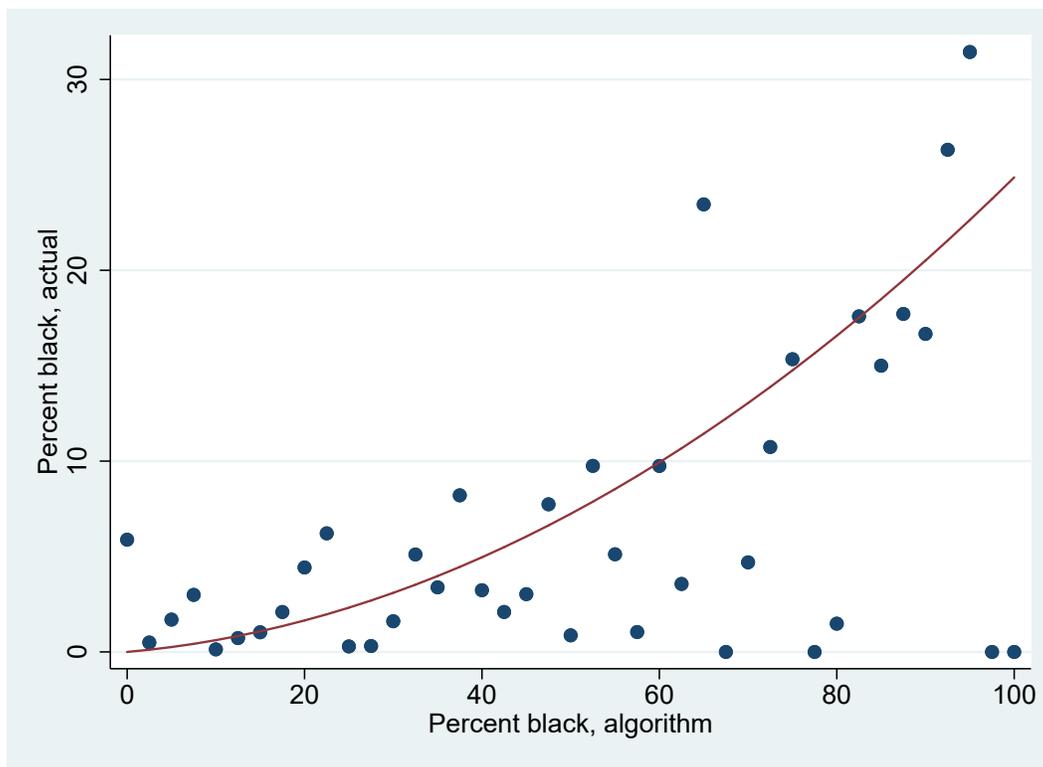
---

<sup>34</sup>The user may indicate that multiple categories of people (friends, spouse/partner, children, other family, coworkers, etc) were present for a given activity.

individuals in these cities, I regress the black interaction rate on CBSA and tract characteristics, and use these relationships (shown in Table 15) to predict the black interaction rate for each city. This exercise suggests that the black interaction rate is slightly higher among Flickr users than among non-black Americans more generally: my estimated black interaction rate for all non-black Americans is 3.5%, compared to 3.7% in the Flickr sample.

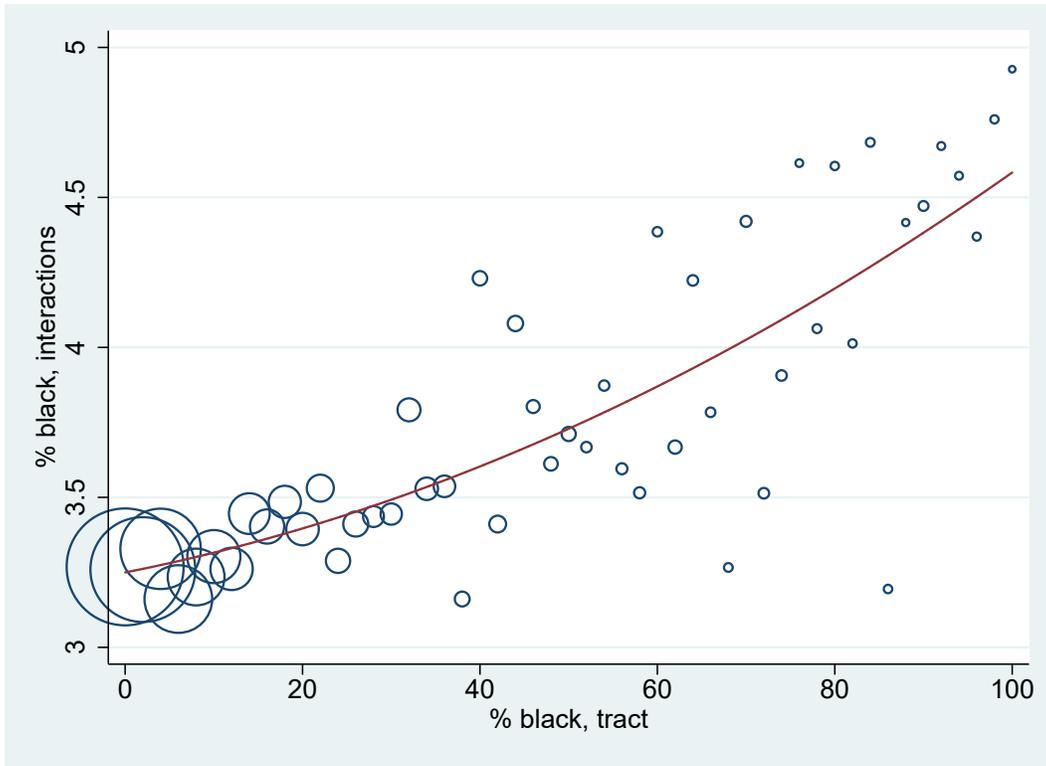
## 10 Appendix Tables and Figures

Figure 4: Relationship between black faces, algorithm vs. hand-coded



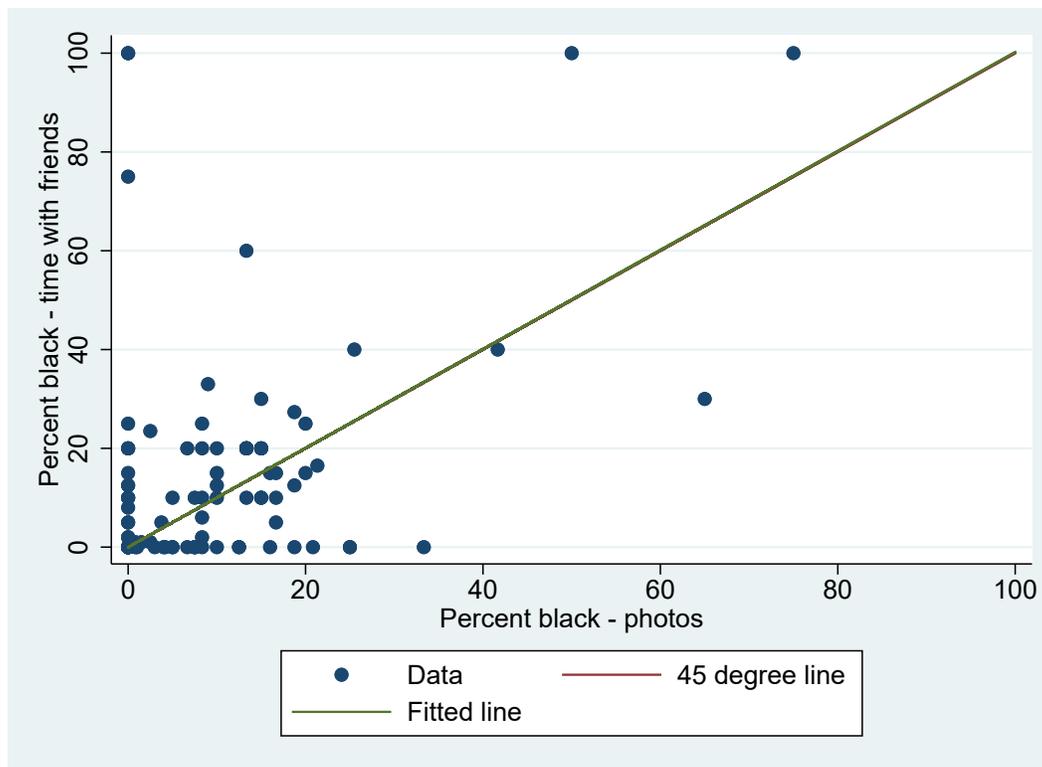
This figure plots the actual fraction of black faces for a set of hand-coded photographs against the fraction of black faces found by the race classification algorithm. Each circle represents a 2.5-percentile group for the fraction of black faces found by the algorithm (0-2.5%, 2.5-5%, etc.); the vertical height shows the mean of the actual fraction black for that group. The size of each circle represents the total number of faces found in the photographs in that group. The red line shows the fitted quadratic relationship between the algorithm's predictions and the actual frequency of black faces.

Figure 5: Relationship between fraction of black interactions and tract demographics



This figure plots the proportion of interaction time spent with black friends against the proportion of the population that is black in a user's assigned home tract. The circles represent the mean within each 2-percentage point cell on the x-axis, with the circle size indicating the number of users in this group. The red line shows the quadratic fit between the two variables. The sample is the set of approximately 88,000 Flickr users living in one of the 148 CBSAs in my main sample.

Figure 6: Relationship between time spent with black friends and black faces in social media photos: MTurk



This figure plots the proportion of interaction time spent with black friends against the proportion of black faces in an MTurk worker’s last social media photograph. The sample is a set of 170 non-black MTurk workers who responded to my survey and spent any time alone with friends on the day prior to the survey. The green line is the fitted line from a regression of the interaction measure on the photo measure (with the constant suppressed); this overlaps almost precisely with the 45 degree line.

Table 7: 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
Coefficient on $\hat{\delta}$	0.094	0.125	0.128*	0.099**
	(0.102)	(0.080)	(0.066)	(0.042)
N	148	148	148	148

This table shows the results from a regression of the mean fraction of photos taken within the indicated distance of Flickr users’ homes on the estimated disutility of travel. A increase in the disutility of travel implies that users living within that city or tract dislike travel more. The sample is the set of 148 CBSAs in my main analysis sample.



Table 8: Frequency of social interactions: ATUS

	Dependent variable: indicator for spent any time with friends on diary day	
	Non-black sample	Black sample
Age	-0.017*** (0.001)	-0.011*** (0.001)
Age squared	0.000*** (0.000)	0.000*** (0.000)
No high school	-0.112*** (0.008)	-0.051** (0.020)
High school	-0.083*** (0.006)	-0.065*** (0.018)
Some college	-0.059*** (0.006)	-0.068** (0.018)
Bachelor's	-0.025*** (0.006)	-0.059*** (0.020)
Region - Northeast	0.004 (0.004)	-0.009 (0.019)
Region - Midwest	0.010* (0.005)	0.002 (0.019)
Region - South	0.002 (0.005)	-0.005 (0.017)
Constant	0.722*** (0.005)	0.596*** (0.038)
N	56,048	9,073
$R^2$	0.028	0.019
Mean of dep. variable	0.245	0.246

This table shows the results from a regression of an indicator for spending any time with friends on demographic characteristics, in the American Time Use Survey sample. The sample is the set of all diary respondents aged 15-85 who filled out a diary on a weekend day.

Table 9: Demographics: MTurk survey respondents compared to U.S. population

	MTurk	U.S. population (over 18)
% age 18-25	15.5%	15.0%
% age 26-45	67.9%	36.0%
% age 46-65	14.5%	33.3%
% age 65	2.1%	15.7%
% male	52.1%	48.5%
% white	73.5%	69.3%
% black	7.1%	11.8%
% Asian	6.5%	5.0%
% Hispanic	11.7%	13.9%
% No high school	0.1%	13.3%
% high school	10.4%	38.0%
% some college	33.3%	23.3%
% Bachelor's	41.3%	16.4%
% post-grad	14.9%	9.0%
N	1,628	11,572,214

This table shows demographic information for my MTurk survey sample, and for the U.S. population over the age of 18. Information on the U.S. population comes from the 2006-2010 American Community Survey.

Table 10: Number of visits to home and other locations in CBSA

	Mean number of visits	Fraction of all visits
Home tract	13.9	41.9%
Other tracts	2.6	4.5%

This table shows the number of unique visits a Flickr user makes to his or her assigned home tract, compared to other tracts she visits. The sample is a set of approximately 84,000 Flickr users who live in one of the 148 CBSAs in my main analysis sample.

Table 11: Number of Foursquare venues around photo locations: home tract vs other visited tracts

	Number of venues		
	Home tract	Other visited tracts	Difference
Food & drink	0.761	0.935	-0.173*** (0.016)
Entertainment	0.546	0.555	-0.009 (0.014)
Stores	0.324	0.450	-0.126*** (0.010)
Offices	0.133	0.154	-0.021*** (0.005)
Other	1.941	2.056	-0.115*** (0.025)
All venues	3.706	4.150	-0.444*** (0.034)

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 is a set of 22,000 Flickr users who live in one of the 148 CBSAs in my main sample. Users were randomly sampled for inclusion in this set, and two photos were randomly sampled for each user: one in the user’s “home” census tract, and one in another tract that the user visits.



Table 12: Relationship between demographics predicted by last name and home tract demographics

	Dependent variable:		
	% same race/eth, CBSA	% same race/eth, visited tracts	% same race/eth, home tract
Prob. user is white	0.073*** (0.004)	0.100*** (0.005)	0.108*** (0.009)
N	9,762	9,762	9,762
$R^2$	0.029	0.034	0.016
Prob. user is black	0.031*** (0.007)	0.053*** (0.009)	0.084*** (0.013)
N	9,762	9,762	9,762
$R^2$	0.002	0.004	0.004
Prob. user is Hispanic	0.109*** (0.004)	0.132*** (0.004)	0.132*** (0.006)
N	9,762	9,762	9,762
$R^2$	0.057	0.093	0.051
Prob. user is Asian	0.044*** (0.003)	0.090*** (0.004)	0.100*** (0.005)
N	9,762	9,762	9,762
$R^2$	0.018	0.058	0.036

The table shows the results from a regression of the proportion white, black, Hispanic or Asian in i) a Flickr user's CBSA, ii) all tracts a Flickr user visits (excluding the home tract), and iii) the home tract on the user's probability of being that race, based on his or her last name. For example, Column (1) in Panel 1 regresses the proportion white in a user's CBSA on the probability that she is white based on her last name. The probability distribution over race by last name comes from the Census Bureau (2010). The sample is the set of Flickr users in one of the 148 CBSAs in my main sample who have a last name appended to their profile.

Table 13: Tract demographics: comparison to U.S. population

	Flickr users	Population - same cities	Population - all
CBSA population	5,552,792	5,006,196	3,834,430
<b>Tract level:</b>			
Density (pop/sq. km)	3,442	2,736	2,197
Median age	37.7	36.9	37.1
Median income	\$34,553	\$30,881	\$29,046
% white	72.0	70.4	73.3
% black	10.4	13.7	12.7
% Asian	8.8	6.1	4.9
% Hispanic	13.4	17.9	16.4
% No high school	11.3	14.7	15.2
% high school	20.1	26.7	28.3
% some college	24.8	27.9	28.2
% Bachelor's	25.5	19.3	17.8
% post-grad	18.3	11.4	10.5
Number of individuals	84,251	216,099,848	285,098,410
Number of tracts	25,661	50,623	66,970
Number of cities	148	148	933

This table shows average home CBSA and tract demographics for users in my sample, compared to the averages for the U.S. population living in the same set of cities (column 2), and to the entire U.S. population (column 3.) Demographic information is taken from the 2006-2010 American Community Surveys.

Table 14: Race classification confusion matrix

	Classified as:	
	Non-black	Black
Actual race: Non-black	87%	13%
Black	25%	75%

This table shows the proportion of non-black/black faces that were classified as non-black/black by the race classification algorithm. The sample is a subset of faces from 20,000 randomly selected Flickr photographs, 10% of which are set aside for testing purposes.

Table 15: Relationship between black interactions and city/tract characteristics: Flickr

	Dependent variable: black interaction rate
Ln CBSA population	-0.005 (0.033)
CBSA segregation	-0.159 (0.487)
CBSA % black	0.008 (0.006)
Ln tract population	-0.097** (0.042)
Tract % black	0.012*** (0.002)
Tract % no high school	-0.002 (0.004)
Tract % high school	-0.001 (0.004)
Tract % some college	-0.011*** (0.004)
Tract % college	-0.006 (0.005)
Tract median age	-0.048 ** (0.024)
Tract median age squared	0.001** (0.000)
Ln tract median income	-0.013 (0.077)
Tract density	0.000*** (0.000)
Northeast	-0.037 (0.100)
Midwest	-0.166 (0.118)
South	-0.155 (0.106)
N	50,360
$R^2$	0.003

This table shows the results from a regression of the black interaction rate for Flickr users on characteristics of their home CBSA and tract. The sample is the set of Flickr users who live in one of the 148 CBSAs in my main analysis sample, and who post any social photographs on weekends.