JTLU

Estimating the effect of land use and transportation planning on travel patterns: Three problems in controlling for residential self-selection

Daniel G. Chatman [a]
University of California, University of California, Berkeley

Abstract: The common understanding of "residential self-selection" generally found in research on the effects of the built environment on travel is in error in three main ways. First, scholars have generally failed to recognize that the built environment may have different effects on travel for different households. Second, controlling for residential self-selection is not necessarily relevant to the predictive questions that controlled estimates are meant to inform. Third, in controlling for preferences and sorting, the literature has failed to account for the composition of the population and its consequences for housing demand. These problems may significantly influence the validity and usefulness of the research.

1   Introduction

A body of research has investigated how the built environment influences travel, in order to find evidence about how transportation and land use planning can influence travel patterns, particularly in order to reduce auto use. The "residential self-selection" strand of this literature has attempted to identify and control for a statistical problem: households may decide where to live based partly on their expected travel patterns, reflecting their heterogeneous and unobserved preferences or needs to carry out activities outside the home. If so, a statistical relationship between the built environment and travel may reflect a correlation with the built environment of those unobserved but heterogeneous travel preferences, rather than an "independent" effect of the built environment upon travel. People who prefer to commute via transit are probably more likely to buy or rent homes near a stop on the rail line to work; those who shop a lot by car may seek a home near a highway on-ramp; people who like walking may seek to live in neighborhoods having good parks within walking distance. Controlling for these ordinarily unobserved and heterogeneous travel preferences—taking them out of the equation, as it were—is thought to yield a better estimate of the built environment's effects on travel, according to this view of the problem.

In this paper I argue that this common conception of residential self-selection is in error in three main ways. Common journalistic explanations of residential self-selection, and even some introductory explanations appearing in the academic literature, do not always betray the various misunderstandings discussed in this paper. But they are found in almost all of the quantitative research itself.

The first problem is that the residential self-selection literature has generally failed to recognize that the "independent treatment effect" of the built environment on travel may *vary* among individuals depending on heterogeneous travel preferences. The idea of an "independent" treatment effect is an assumption, not a fact. The second problem is that controlling for residential self-selection is not necessarily relevant to the predictive questions to which controlled estimates are intended to answer. That is, residential sorting is not a control problem—it's the "name of the game." And third, in controlling for preferences and sorting, the literature has forgotten to look at the question of the composition of the population and consequences for housing demand, an issue that is even more important if different groups are more or less responsive to changes in the built environment. For these reasons, much of the research controlling for self-selection may fail to inform the research question that motivates it.

---

[a] dgc@berkeley.edu

2   The existence of an "average" built environment effect is only an assumption

> With so much attention on explaining behavior in terms either of na-
> ture or nurture, scientists at the University of California, San Francisco
> recently described a fascinating example of how heredity and environ-
> ment can interact. Perfect pitch is the ability to recognize the absolute
> pitch of a musical tone without any reference note. People with perfect
> pitch often have relatives with the same gift, and recent studies show
> that perfect pitch is a highly inherited trait, quite possibly the result of
> a single gene... But the studies also demonstrate a requirement for ear-
> ly musical training (before age six) in order to manifest perfect pitch.
> Time will tell whether there is a "perfect pitch" gene, but it seems rea-
> sonable to think that many personality and behavioral traits will not be
> exclusively the province of nature or nurture, but rather an inextricable
> combination of both. (Davies 2001)

> Explanatory variables have been treated as independent when, in fact,
> they almost certainly have combined effects on transit ridership that
> are more or less than the sum of the parts. For example, improvements
> in transit service may have little effect in automobile-rich areas but a
> large effect in automobile-poor ones. (Messenger and Ewing 1996)

Just as nature and nurture interact, and as built environment characteristics themselves may have combined effects that exceed their individual effects, so it can be said that preferences may influence travel behavior in combination with built environment characteristics, and not entirely apart from them. As "nature" and "nurture" cannot be independently conceived, neither can preferences and the built environment be assumed to be independent of each other.

The built environment "treatment" upon travel choices can be thought of as a set of prices of traveling for different modes (Boarnet and Crane 2001). Dense, mixed-use places may have lower prices of walking and higher prices of driving, in time terms. Such neighborhoods would likely induce walking in any random person dropped into that neighborhood who had formerly been randomly assigned to a neighborhood where by virtue of the built environment, the price of driving was low, and the price of walking was high. The premise of the residential self-selection research is that people choose where to live based on seeking lower prices for their desired travel patterns—in the form of shorter travel times, more convenience and safety, and lower financial costs.

Of course, the price elasticity of demand can also vary among individuals. In fact, one definition of having a "preference" for something is having relative price inelasticity of demand for it. For example, a person with a strong transit preference might be less dissuaded from using transit by the removal of a nearby stop, or by a decrease in bus speeds—even though the same person will be more personally affected by such a change. By the same logic, the amount of ice cream consumed by a person with a strong preference for ice cream will be less affected by an increase in the price of ice cream—even if he is more annoyed by a price increase.

Whether there is very much preference-driven variance in travel demand with respect to the built environment is, of course, an empirical question. People with stronger preferences for a particular travel mode may in some circumstances be more affected by a change in the built environment. It likely depends on the range of price changes. For example, it is quite difficult to use transit if there is no transit stop within a reasonable distance of the places one is traveling from, or near the places one wants to go. There may be a particularly big jump in transit use among those with strong transit preferences once transit service is made available. But larger improvements, or improvements in some other range of effec-

tive prices, might have little effect on the use of transit by those with a strong preference to use transit. Meanwhile, large improvements might begin to influence the transit use of those who do *not* prefer transit.

The differential responsiveness to changes in the built environment of people with different travel preferences would not much matter, except that the entire premise of the residential self-selection hypothesis is that people decide where to live based on those very preferences. It is helpful to visualize the different assumptions embedded in the question of whether there is differential responsiveness to the built environment, in the context of a cross-sectional study in which residential sorting has taken place. I illustrate using the frequency of walking and its relationship to a built environment index of "walkability" (comprised by weighted measures of intersection density, sidewalk connectivity, mixed uses, and other pedestrian amenities). Let us assume for the sake of illustration that in neighborhoods with the lowest walkability, only people without a strong preference for walking have chosen to live (Group A), and at the highest level of walkability, only people with a strong preference for walking have chosen to live (Group B). Meanwhile, in the range of neighborhoods with intermediate levels of walkability, there is a linear increase or decrease respectively in the share of the population in the two groups. Lastly, for the sake of these illustrations, let us assume that the total population is evenly split between groups A and B.

Figures 1 and 2 illustrate the conventional understanding of the residential self-selection problem. The black lines represent the actual levels of walking as they vary with walkability; the red line is what we observe if preferences are unknown. In both Figures 1 and 2, preferences and the built environment do not interact with each other. The problem of controlling for residential self-selection is limited to estimating (that is, controlling for) the separate fixed effect of preferences for walking.

In the world depicted in Figure 1, the built environment does not influence travel. It is simply correlated with travel preferences because of sorting. A model ignoring those preferences would erroneously conclude that the built environment does influence travel when it does not—because in this world, the built environment influences only where people with different travel preferences choose to live.
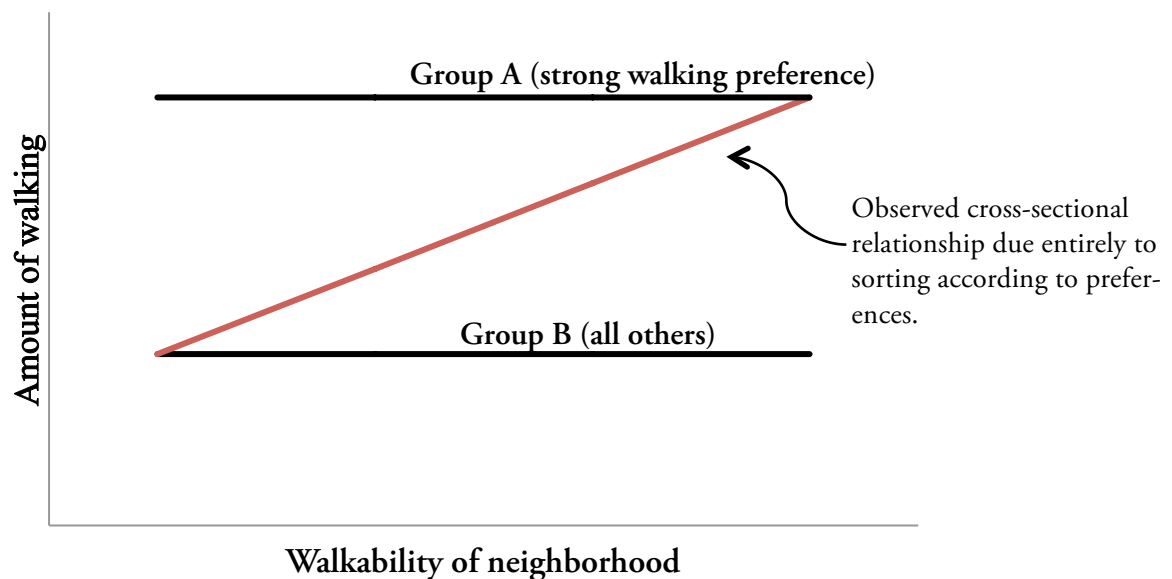

Figure 1: Walking preferences have a fixed effect on walking; the built environment has no influence.

In the world depicted in Figure 2 (below), the built environment does influence travel, and by the same amount for everyone. There is, again, a fixed effect of preferences upon walking. But the influence of walkability is much less than might be inferred from a statistical model that failed to control for the fixed effects of walking preferences. Observing walking preferences, and controlling for them, corrects for the false inference in this case.
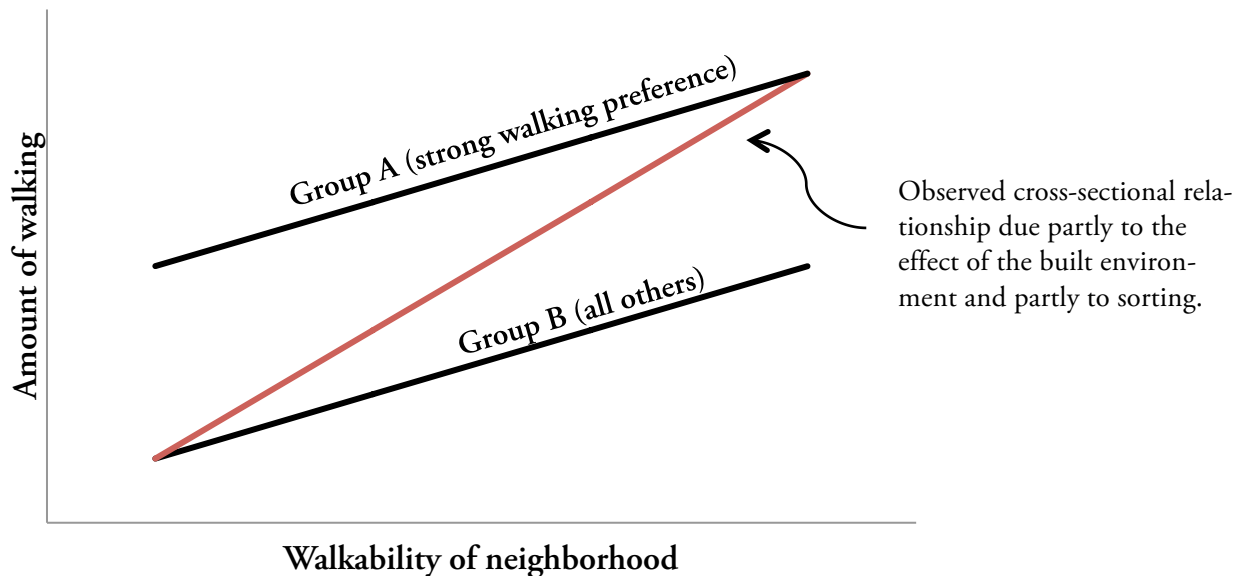

Figure 2: Walking preferences have a fixed effect on walking; the built environment influences walking the same way regardless of preferences

In the scenarios shown in Figures 1 and 2, controlling for walking preferences in the conventional way would yield correct estimates of the built environment's effect on travel. But the facts of the matter may not be like those depicted in Figures 1 and 2. Scholars generally have not been paying any attention to the possibility that preferences might be associated with different levels of response to the built environment, in addition to influencing residential sorting.

Figures 3 and 4 illustrate two (of many) possibilities along these lines. In Figure 3, the level of walking is equal among those with strong and weak preferences in the lowest-walkability neighborhoods, and those with strong preferences are more responsive to increases in walkability. In Figure 4, people with strong preferences for walking have a higher base level of walking regardless of walkability; and they respond less to increases in walkability than do those without a strong preference for walking.

In either case, the net effect of a built environment change depends on sorting according to travel preferences, and also upon the share of the population with different preferences—precisely because of the fact that in these scenarios, there is differential responsiveness according to the same travel preferences that influence residential location. The premise of correcting for residential self-selection is that sorting confounds our estimates of the effect of built environment changes. But if responsiveness varies according to preferences (e.g., if those with a strong preference to walk respond less to sidewalk availability than the rest of the population), then the "average" or "independent" built environment treatment effect is the wrong answer, precisely because of the fact that there is sorting by preferences.

The built environment effect in the scenarios depicted in Figures 1 and 2 does not vary according to the composition of the population or how much sorting there is. In the scenarios depicted in Figures 3

and 4, where there is differential responsiveness to the built environment, these facts do indeed matter in calculating any treatment effect.
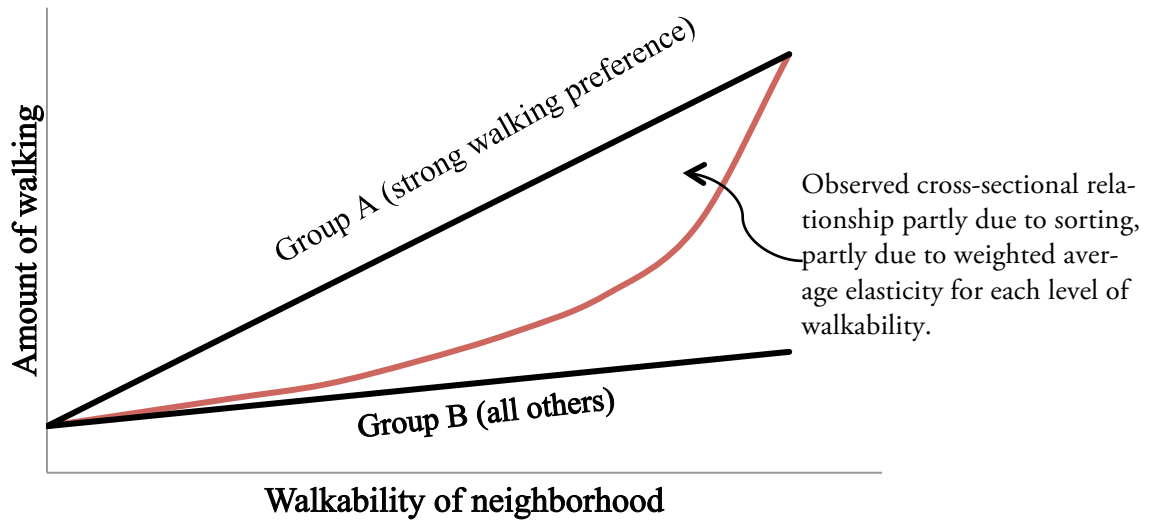

Figure 3: People with walking preferences are more responsive to the built environment.
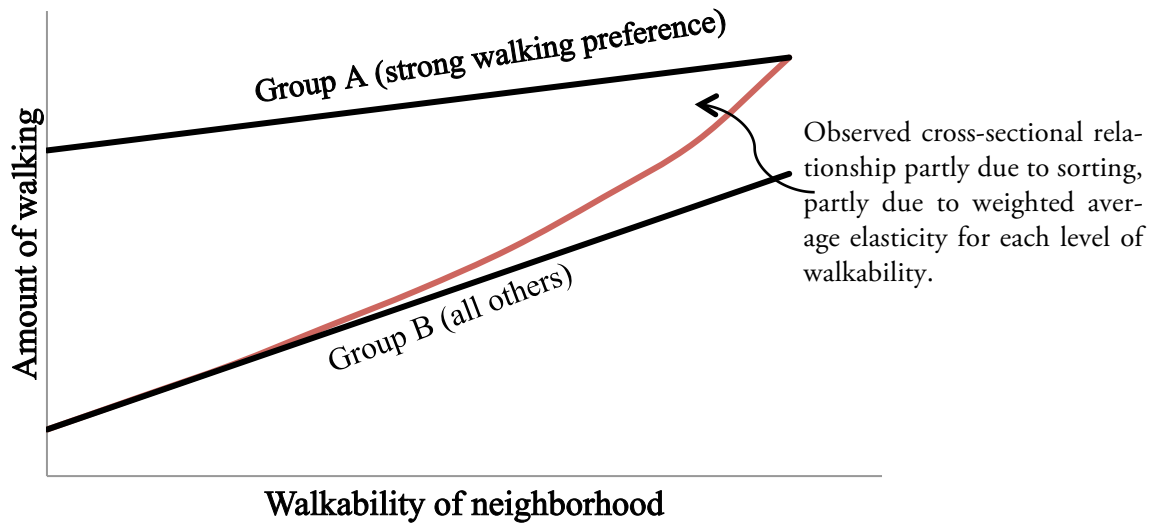

Figure 4: People with walking preferences are less responsive to the built environment, but have a higher rate of walking.

There is a potentially important and related misconception, almost universal in this literature, that treatment effects of the built environment will be overestimated among the general population if not controlling for residential self-selection (with Chatman 2009 and Naess 2009 providing two exceptions). I speculate that this means most scholars believe two things: (a) self-selecters for a particular mode of travel—those with strong preferences for that mode—always use that mode more than people who don't self-select *ceteris paribus*; (b) the effect of the variation in the built environment upon travel is about the

same for everyone regardless of preferences. The former assumption seems reasonable enough. The latter assumption is an empirical question but on first blush seems rather unreasonable. If the world is like that depicted in Figures 3 or 4—that is, if responsiveness to the built environment varies according to travel preferences—then the common technique to correct for residential self-selection is flawed. The treatment effect on a randomly selected group is the wrong research question entirely. The correct question is: What will be the treatment effect on the group that will experience the change in the built environment?

The notion of differential travel-behavior responsiveness to the built environment with respect to demographic characteristics has been present for a long while (e.g., Messenger and Ewing 1996; Badoe and Miller 2000) but it has not necessarily been extended to the notion of differential responsiveness with respect to ordinarily unobserved heterogeneous preferences. Bhat and Guo were perhaps the first to operationalize the interaction in a joint sense with residential choice, which is an extraordinarily important and underappreciated contribution—though they did so only with respect to demographic characteristics, particularly income, and not with respect to preferences for auto use (Bhat and Guo 2007). In general, even the most sophisticated studies methodologically (e.g., Pinjari, Bhat, and Hensher 2009) or those proposing to use the best longitudinal panel datasets (e.g., Boone-Heinonen et al. 2011) fail to recognize the possibility *methodologically* that BE has differential effects based on preferences, even if they may identify that possibility conceptually.

This omission might be a grave one in understanding the effects of residential self-selection. We cannot calculate the correct treatment effect of the built environment by treating preferences as having fixed effects on travel if they interact with the built environment. How far off our estimates will be depends not just on residential sorting but also on the difference in the built environment effects between those with different preferences, and on the share of the population accounted for by people with different preferences. I turn to those issues next.

3   Residential sorting is the name of the game

> Ignoring self-selection with respect to mode preferences could easily lead to an overestimation of changes in mode choices after changes in level-of-service characteristics. For example, if a new railway station is built in an existing neighbourhood, the share of people with a preference to travel by train in that neighbourhood will probably be lower than average. In this case, current models might overestimate the share of rail. (van Wee 2009)

When moving to a new neighborhood, some people likely choose where to live in order to be able to travel in a way that they could not in their old neighborhood. Other people choose particular neighborhoods so that it is easier to travel the way they are already traveling. And still other people will choose neighborhoods without focusing very much on how the choice will affect their ease of travel by different modes, since many other important factors play a role in deciding where to live.

A model controlling for residential self-selection, though estimating correct treatment effects on a randomly chosen population group, could actually overestimate the causal effect of a rail investment. If differential elasticities of demand exist with respect to rail access, then the estimates would be off even further. In this case one needs to make more explicit assumptions about the composition of the group in order to correctly estimate travel patterns of residents already living in the area.

The epigraph assumes that preferences for rail are associated with greater responsiveness to rail access. While this shows an awareness of the possibility of differential responsiveness, the assumption could be wrong. It is better understood as an empirical question that has been largely ignored in empirical study. If one wants to know the immediate effect of the opening of a rail station upon travel patterns, and if responsiveness to transit access depends on preferences and demographic characteristics, then one needs

to know something about both the differential effects of transit access for different groups, and the composition of the population living and working near the rail station. This has simply not been considered when using the "controlled" built environment effect to predict the travel impacts of built environment changes.

If the research question is "what is the independent effect of the built environment on travel," then interacting attitude measures with built environment effects, or allowing flexibility of parameter estimates among groups identified as having such differential preferences, will yield a more accurate but more complex answer—not just an average treatment effect, but a group-specific treatment effect; or an interactive effect between preferences and the built environment. It will be a complex answer—one that will vary depending on assumptions about the distribution in the population of heterogeneous preferences, and upon assumptions about or estimates of heterogeneous elasticities of demand with respect to the built environment.

However, if the research question is, for example, "What happens when a dense housing development is built near a rail station?" then we have a different problem. The question of an "effect" from a regional perspective is a net zero if the world is like Figure 1. In that case, transit-preferring people would always use transit more than the rest of the population and the built environment has no direct influence upon transit use. If the world is like Figure 2, then the answer is just the effect due to however many more households now live near transit than was the case before. But if responsiveness to transit access varies according to transit preferences—that is, if the world is more like Figures 3 or 4—then there is no such thing as an independent built environment treatment effect. One needs an answer that takes into account sorting. The causal question is to understand how a change in development (e.g., new housing near a rail stop) causes both residential choices and travel changes.

When we control for residential self-selection, we are often, perhaps usually, weeding out the important aspect of residential choice in answering research questions. This is not a new observation; Jonathan Levine was one early critic of residential self-selection research because he saw residential choice as in fact the important issue in providing alternative housing development (Levine 1998). But the observation bears repeating, because it becomes more important if differential responsiveness to the built environment is associated with demographic characteristics, attitudes, and preferences.

In some cases, as with a new housing development, the entire premise of a change in the built environment is in fact an influx of population. Residential choice is the name of the game, and controlling for residential self-selection means potentially missing out entirely as long as the influx of population is not a randomly selected group—and of course it will not be a randomly selected group, stipulating that one needs to control for residential self-selection in the first place.

Therefore the fact that residential self-selection is likely going on, and that there may be a difference in responsiveness of the different kinds of households making different residential location decisions, implies a need to consider both "uncontrolled" and "controlled" analysis. For example, for analysis that is meant to predict what will happen if housing is developed near rail stations region-wide, we might expect sorting by observed demographic characteristics to explain some of the observed differences in travel patterns. But these changes are really to be thought of as part of the *causal effect* of the change in the built environment (Cao and Chatman 2011). We might even be better off with uncontrolled estimates in many cases. How do we know that the controlled ones are better, if net travel pattern impacts depend on sorting?

4   Policy relevance requires a larger supply and demand perspective

To some extent, "controlling for preferences" means throwing away information, including information about the distribution of preferences in the population for different built environments; the likely residential sorting responses to BE changes; and the potential demand for BE changes given heterogeneous preferences. In other words, preferences and their distribution in the population have the potential to

provide valuable information about markets for alternative development, and therefore to predict travel patterns. This information may even be more valuable than estimates of average effects of the built environment on travel.

Residential choice in the context of built environment policies that are intended to reduce auto use implies a great deal of complexity indeed. For example, zoning reform may lead to a mixed-use neighborhood; a person who prefers to walk may relocate into that neighborhood; and that person may walk more to shops than she did in the previous neighborhood, both because of her pre-existing preferences and because the built environment makes walking easier or more pleasant (Figure 5, below). But as a number of scholars have pointed out, the causal influences work in all directions. Moving into a neighborhood may affect her preferences or attitudes, because her options have changed. The increase in habitual walking could also influence her preferences. And her habits themselves, entirely separate from her preferences, may lead to the choice as being seen as acceptable even if perhaps her preference would be to walk less.
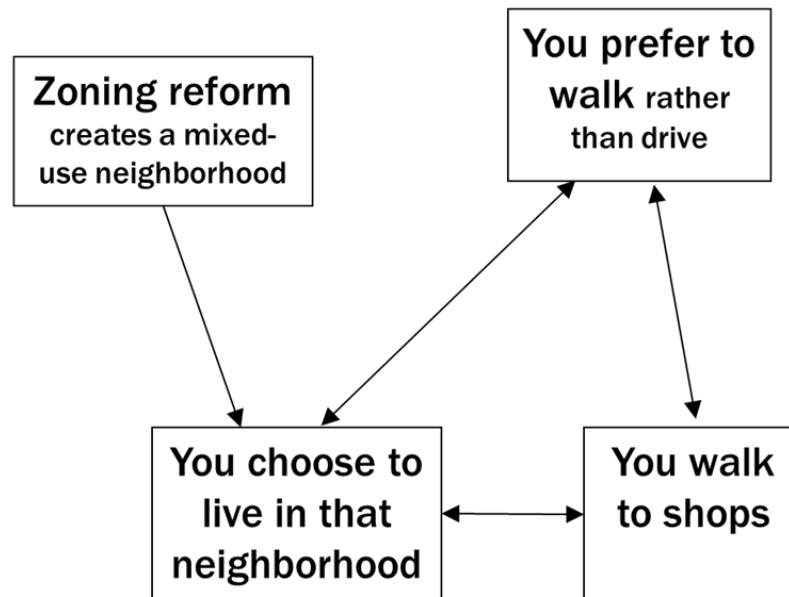

Figure 5: Residential self-selection in the context of a policy change

It gets even more complex. In Figure 5, there is only one unidirectional arrow, that drawn between the zoning change and the choice to live in the neighborhood. But what leads to the zoning reform? No individual's demand does so—but aggregate demand can do so. For example, developers may pressure planning agencies to permit mixed use development in response to their understanding about market demand for such development.

Research that controls for residential choice does not really begin to touch questions about market demand, and it cannot if it does not take account of the possibility that there are heterogeneous travel responses to the built environment associated with different market segments for housing. It is not clear, in fact, that the "independent built environment effect" is nearly as relevant as other questions, if one wants to know what will happen when a policy change is adopted. The answer depends on the regional supply and demand for housing, which suggests a strong need to understand what happens to settlement patterns when investments in transport or land use policy changes are carried out.

5   Conclusions

There are three interrelated problems neglected by conventional research controlling for residential self-selection. The first is that there may be differential responses to the built environment by different groups of people, whether or not those differences are directly observable (like household income) or not (like mental states such as travel preferences). If so, there is actually no such thing as an independent treatment effect of the built environment. To put it another way, the average treatment effect is not relevant, since the extent of responsiveness varies according to preferences, the built environment, and the combination of the two. The most methodologically sophisticated approaches do not make this problem go away; it needs to be woven into those approaches. There will be challenges in doing so because of the fact that differential elasticities with respect to preferences are difficult to model, since preferences are not directly observable and are themselves influenced by multiple factors: the built environment, habitual travel patterns (because of experience effects), and cognitive dissonance effects (Chatman 2009).

The second problem, exacerbated by the first, is that BE-travel relationships actually consist partly of sorting according to heterogeneous preferences. And the third problem, one that potentially dwarfs the others, is that some policies might have big effects on travel and others might not, depending on the population composition of the market (that is, the distribution of preferences in the population).

The potential existence of heterogeneous elasticities with respect to the built environment also means the need for more empirical studies of residential choice and a better understanding of the share of metropolitan area populations seeking and not finding housing that meets their travel preferences (Cao and Chatman 2011). For example, if most of the environmental benefits of new built environment policies are due to the relocation of households with strong preferences for alternative modes, then estimating the net regional benefits of such a policy depends on a better understanding of the elasticity of demand with respect to the built environment of travel for that mode, by people in that preference group—not for an average effect across the population, even if the average effect could be correctly estimated.

Finally, these problems with the way residential self-selection has been carried out also have broader implications for research that controls for demographic characteristics. Just as controlling for preferences renders estimates of built environment effects potentially wrong, the effect of demographic characteristics on both built environment effects, and upon sorting, is ignored in conventional analysis. Demographic factors may mediate the influences of the built environment on travel, and they certainly influence residential choices. The built environment may have larger effects for poorer people, for immigrants, and so on. Or they may have smaller effects. We don't know, because we generally have not tested this, beyond the rare exceptions discussed earlier. When we universally control for demographic characteristics in estimates of the built environment on travel, we neglect the fact that those demographic characteristics are also associated with residential choice, and that their travel demand may vary along with those demographic characteristics.

References

Badoe, D. A., and E. J. Miller. 2000. Transportation–land-use interaction: Empirical findings in North America, and their implications for modeling. *Transportation Research D* 5 (4): [235]-263.

Boarnet, M. Gary, and R. Crane. 2001. The demand for travel. In Travel by design: the influence of urban form on travel. Oxford; New York: Oxford University Press.

Boone-Heinonen, J., P. Gordon-Larsen, D. K. Guilkey, D. R. Jacobs Jr, and B. M. Popkin. 2011. Environment and physical activity dynamics: The role of residential self-selection. *Psychology of Sport and Exercise* 12(1):54-60.

Cao, X., and D. G. Chatman. 2011. How will land use policies affect travel? The importance of residential sorting. Transportation Research Board 91st Annual Meeting.

Chatman, D. G. 2009. Residential choice, the built environment, and non-work travel: Evidence using new data and methods. *Environment and Planning* A 41(5): 1072-1089.

Davies, K. 2001. Nature vs. nurture revisited. NOVA.

Levine, J. 1998. Rethinking accessibility and jobs-housing balance. *Journal of the American Planning Association* 64(2): 133-150.

Messenger, T., and R. Ewing. 1996. Transit-oriented development in the Sun Belt. *Transportation Research Record* 1552: 145-153.

Pinjari, A. R., C. R. Bhat, and D. A. Hensher. 2009. Residential self-selection effects in an activity time-use behavior model. *Transportation Research Part B-Methodological* 43(7): 729-748.

van Wee, B. 2009. Self-selection: A key to a better understanding of location choices, travel behaviour and transport externalities? *Transport Reviews* 29(3): 279-292.