Title

Estimating the Travel Behavior Effects of Technological Innovations from Cross-Sectional Observed Data: Applications to Carsharing and Telecommuting

Author
Mishra, Gouri Shankar

Publication Date
2017

Estimating the travel behavior effects of technological innovations from cross-sectional observed data: Applications to carsharing and telecommuting

By

GOURI SHANKAR MISHRA

M.S. Transportation Technology & Policy (University of California, Davis) 2010
M.S. Finance (Birla Institute of Technology & Science, Pilani) 1998
Master of Management Studies (Birla Institute of Technology & Science, Pilani) 1998

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in
TRANSPORTATION TECHNOLOGY & POLICY

in the
Office of Graduate Studies

of the
UNIVERSITY OF CALIFORNIA
DAVIS

Approved:

____
Chair: Dr. Patricia L. Mokhtarian

____
Dr. Sonia Yeh

____
Dr. Keith F. Widaman

Committee in Charge
2017

ABSTRACT

In this dissertation, we estimate effects on travel behavior of two specific technological innovations – emerging shared mobility services and telecommuting – using publicly available travel surveys. These surveys are cross-sectional and observational in nature, which leads to the potential for (1) selection bias due to observed and unobserved differences in characteristics between program participants and non-participants; and (2) reverse causality bias arising because of potential influence of the travel behavior outcome of interest on the propensity to enroll in the program. Our methodological framework combines established methods from both statistical and econometric literature to draw causal inferences. The key innovations in this dissertation are the combination of diverse methods to address the joint occurrence of various biases, and their specific empirical applications. We also compare the results of alternative methods.

In the first study (Parts II & III of the dissertation), we estimate the effect of carsharing on travel behavior, using data on employed San Francisco Bay Area respondents from the 2011-12 California Household Travel Survey (CHTS). We find that 80% of the observed difference of 0.9 units in average vehicle holdings between carsharing non-members and members may be explained by self-selection and reverse causality biases. The remaining difference of 0.17 units reflects the estimated effect of carsharing, which is the equivalent of shedding one vehicle by about one out of every six households whose member(s) are enrolled in carsharing. The effect on transit usage and walking and biking frequency is positive, albeit small and statistically non-significant.

In the second study, we estimate the effect of the adoption of telecommuting on travel behavior for full-time employed respondents with a fixed work location outside home, using data from the annual United Kingdom National Travel Surveys for the years 2009 to 2013. On average, telecommuters are observed to travel more than non-telecommuters. However, after accounting for the *observed* differences in *traits* and *tastes* between the two groups using a linear

regression model, the differences fade to (nearly-) insignificant levels. Further control of self-selection bias arising from *unmeasured* differences in "relevant" characteristics leads to the conclusion that telecommuting has a substitution effect on both commute and non-work travel.

Our results are broadly consistent with those of earlier studies, which, unlike our study, are based on purpose-built proprietary surveys explicitly designed to evaluate effects of either of the two programs. Although the data collected through those other means are still observational in nature, various biases identified in this dissertation may be addressed by questionnaire design, including retrospective reporting of travel behavior before and after enrollment in the program. By implicitly assuming that the unobserved influencers of both program adoption (either telecommuting or carsharing as the case may be) as well as travel behavior do not change over the course of the evaluation (an assumption which may or may not be true), those prior studies estimate effect by measuring change in travel behavior before and after program enrollment relative to a control group. Unfortunately, such surveys are expensive, proprietary, and usually one-off studies.

Large regional travel surveys, on the other hand, are publicly available, leading to the potential for replicability and involvement of multiple research teams. Further, these surveys collect information about broader travel behavior patterns and yield samples that are often larger and more representative of the general population. However, the cross-sectional and observational nature of these surveys creates the potential for joint occurrences of various biases identified in this study, which makes it necessary to adopt methodologies that correct and control for these biases when estimating causal effects. We hope that the methodological frameworks adopted in our study will provide an example that other researchers can use to analyze various programs in transportation using publicly available travel surveys, and that the causal inferences drawn will offer a sound basis for policymaking.

CONTENTS

---

[i] Paper published in Research in Transportation Economics 52 (2015): 46-55. DOI: 10.1016/j.retrec.2015.10.010
[ii] Paper resubmitted for review to Transportation (Springer Journal, January 2017)

[iii] Manuscript in preparation

ACKNOWLEDGEMENTS

I have many people to thank for making this dissertation possible and helping me in my PhD journey.

Without question, I should start with my advisors. Patricia Mokhtarian graciously accepted to first work with me to get me started in the disaggregate travel behavior area when I re-admitted to UC Davis for my Doctorate in 2013, and then subsequently to be my primary advisor when I decided to focus my dissertation in that area. She has patiently mentored me in the multiple technical areas necessary to complete my dissertation – but in the process ingrained in me the value of empathy and appreciation of diverse perspectives.

I came to UC Davis first as a Master student in 2009 under the guidance of Sonia Yeh who introduced me to the world of energy, climate change, and environmental modeling. We continued to work together in those areas during my PhD even though they do not form part of my dissertation. I have had the good fortune to have Keith Widaman (now at UC Riverside) as an advisor since I took a series of classes in Psychometrics from him. Although Lew Fulton and Richard Plevin were not my "official" advisors, they always played the role. I worked on multiple projects with both of them and these were great learning experiences for me. Rich encouraged me to take a counterview to environmental modeling, and that exercise was very gratifying and stimulating; and also influenced my research for the dissertation.

All of my advisors - official and unofficial - have provided me with excellent advice and feedback, constantly encouraged me to broaden my perspective and improve the quality of my work, accommodated my personal constraints & quirks, and almost always uplifted my spirits during our interactions.

Ken Laberteaux from Toyota and I had extensive discussions on the effect of telecommuting on travel behavior when I was advising Toyota on the topic prior to my enrollment at UC Davis for PhD. Subsequently, we also had extensive discussions on the topic of emerging mobility – carsharing and on-demand ridesharing. Both topics ended up being the core of my dissertation work and I greatly benefited from our earlier discussions. While working on the second topic, Ken introduced me to Regina Clewlow, first at Stanford and now at UC Davis. I ended up working closely with Regina on carsharing and we co-authored two papers that form part of my dissertation.

I would like to thank my fellow doctoral students - especially Jacob Teter and Geoff Morrison for their feedback and advice on academic and professional matters, and patience during our joint research work. We worked together, along with Page Kyle from Pacific Northwest National Laboratory, on a number of projects and co-authored papers. I had a number of stimulating discussions with all my co-authors and learned greatly from them.

Several other faculty members and researchers greatly contributed to this dissertation. Susan Handy, Gil Tal, David Bunch and Joan Ogden, who were also part of my Qualifying Exam Committee, helped me frame my dissertation objectives and in the process influenced the research direction. I also benefited from my long discussions with Rosa Dominguez-Faus on topics related to energy and environmental modeling.

I also want to acknowledge the support from the support staff of Institute of Transportation Studies. Specifically Annemarie Schaaf, the Graduate Program Coordinator - who went out of her way to resolve administrative and financial issues faced by the students including me. Ernie Hoftyzer ensured that I was able to attend all Asilomar Conferences during my time at UC Davis – these conferences on transportation and energy policy were an excellent opportunity to learn and get feedback from industry leaders. Others who supported me in my journey include Helen Barr, Beth Bourne, Ning Wan, Valerie Johnson and Natalie Ruiz.

Of course, this dissertation would also not have been possible without the financial support of several organizations. The Graduate Scholars Fellowship and Graduate Research Mentorship Fellowship both from University of California Davis; the National Center for Sustainable Transportation (NCST) Dissertation Fellowship; the Henry A. Jastro Endowment;

and the BP Corporate Affiliate Fellowship. I also received funding from Sustainable Transportation Energy Pathways (STEPS).

There are a number of other individuals who influenced my PhD journey. Jim Sweeney of Stanford University provided me an opportunity to organize the weekly Sustainable Mobility Seminar series at Stanford and learn from various experts during that period. Than Seeds and Donald McKay encouraged me to pursue graduate studies when I worked with them at APL prior to UC Davis.

The biggest gratitude I have is towards my wife and parents. My decision to stop working (and hence earning) and pursue a PhD had opportunity costs. My decision to pursue research in two distinct fields - energy and environmental modeling, and travel behavior modeling – amplified the required time and hence opportunity costs. My wife - Neelam - supported this decision and was the sole breadwinner for all these years. My parents and brother also encouraged me to pursue the Doctorate studies with promise of a "financial insurance". This journey would not have been possible without the support of my family.

PART-I

Summary

SUMMARY

1. Background and Dissertation objectives

It is often stated that the gold standard for estimating the effects of a treatment or program is a longitudinal study with randomized assignment of respondents to the program or no-program (control) groups (Rubin 2008). Random assignment reduces the chance of systematic differences between the treatment and control groups before program assignment and thus (to the extent that it does result in statistically equivalent groups) controls for selection bias. Longitudinal data confirm the temporal sequence of events – outcome (usually) follows the treatment, allowing analysts to distinguish the effect of program membership on outcome from the effect of outcome status on propensity to join (or remain in) the program. Unfortunately, many evaluations are based on cross-sectional observational studies, and therefore possess the potential for both selection bias and simultaneity bias.

In this dissertation, we estimate effects on travel behavior of two specific technological innovations – emerging shared mobility and telecommuting – using publicly available travel surveys, which are both cross-sectional and observational in nature. There is a rich literature on estimating the effects of both innovations on travel behavior, but in most cases the estimates are based on purpose-built surveys. Although the data collected through these means are still observational in nature, some of the biases identified in this dissertation can be addressed by intelligent questionnaire design, especially retrospective reporting of travel behavior before and after enrollment in the program. By implicitly assuming that the unobserved influencers of both program adoption (either telecommuting or carsharing as the case may be) as well as travel behavior do not change over the course of the evaluation (an assumption which may or may not be true), the studies estimate effect by measuring change in travel behavior before and after program enrollment relative to a control group. Unfortunately, such surveys are expensive and often proprietary, and usually one-off studies.

A key goal of this dissertation is to capitalize on publicly available household travel surveys like the nationwide United Kingdom National Travel Survey (UK-NTS) and regional California Household Travel Survey (CHTS). These surveys usually yield large samples that are representative of the general population, collect detailed information about demographics and travel behavior patterns, and most importantly, are publicly available. On the downside, these surveys do not capture the dynamic aspects of travel behavior. Specifically, it is not possible to confirm the temporal sequence of events, leading to the potential for reverse causality. Further, these surveys usually do not collect information to gauge respondents' attitudes and preferences towards such issues as environment, driving, and technology, which are known to confound the effect estimates.

We estimate effects using well-established methods from the statistical and econometric literature. The key innovation in this dissertation is to combine diverse methods to address the joint occurrence of various biases. We also compare the results of alternative methods. In Part-II we use a combination of propensity score and regression methods to estimate the effects of enrolling in carsharing on vehicle holdings, frequency of transit usage, and frequency of biking and walking, but conclude that the effect estimates are unreliable because not all biases have been addressed. We extend the research in Part-III and use Structural Equation Modeling to control for the remaining biases and estimate the effects of non-motorized and transit trips. In Part-IV, we use two procedures widely deployed in the econometric literature to estimate the effects of telecommuting. Both these procedures use instrumental variables and are called the two-stage predictor substitution (a close relative of two-stage least squares, or 2SLS) and Control Function (specifically the Heckman Treatment Effect Model) approaches. Part-II reproduces our manuscript published in the journal *Research in Transportation Economics*. Part-III reproduces our manuscript resubmitted to the journal *Transportation* and currently under review (February 2017). Part-IV is being finalized for submission.

In this summary to the dissertation, we describe the biases that need to be addressed, and identify methods from statistical and econometric literature to control for these biases. Finally, we briefly recapitulate the results of our analysis and also our experiences with the application of these methods to estimate the effects of emerging shared mobility and telecommuting on travel behavior.

2. Potential biases

In Table 1, we summarize the five key biases that could potentially arise when making causal inferences using cross-sectional observed data.

Observational study participants may have selected themselves into the program under consideration – hence naïve comparison of outcomes between treated and control members reflects not only the effect of the program under study, but also selection bias. Selection bias may arise from differences in characteristics between treatment and control members. The characteristics of interest are those that influence the decision to join the program, as well as the behavioral response to the program. These characteristics could include individual and household demographics, attitudes towards issues (such as the environment or adoption of new technology), employer policies and social norms. When these characteristics are observed, selection bias may arise in two possible ways (Heckman et al. 1996). Selection bias may arise due to non-overlapping observed characteristics (B1; see Table 1), i.e. observations *outside* the common support, in the parlance of the treatment effects literature. In other words, this bias occurs when for some control members there are no comparable program participants with similar characteristics, and/or vice versa. Additionally, selection bias results from differences in *distribution* of observed characteristics *within* the common support (B2). Differences in distribution refer to differences in expectation (average value), in dispersion (usually referring to variance or its square root, standard deviation), and/or in shape (usually summarized by third and fourth moments, i.e. skewness and kurtosis).

If the observed differences are assumed to entirely account for the selection bias, then the situation is variously referred to in the literature as unconfoundedness, selection on observables, exogeneity, and conditional independence (Elwert and Winship 2014). A more reasonable assumption is that not all variables affecting program participation or selection are observed by analysts or researchers, yielding selection on unobservables. Such a situation leads to an omitted variable bias (B3) (Angrist and Pischke 2008 pp. 44-47; Elwert and Winship 2014; Heckman et al. 1996).

In addition to self-selection bias, studies estimating causal effects of programs from cross-sectional data may have to confront the potential for reverse causality bias or simultaneity bias (B4). Simultaneity bias is produced when an "explanatory" variable is simultaneously a function of the "dependent" variable it is supposed to explain – that is, when one variable is both a cause and an effect of another (Mokhtarian and Cao 2008).

Table 1: Potential biases while drawing causal inferences from cross-sectional observed data, using examples from carsharing based on the 2011-12 CHTS dataset.

Bias	Source	Example: Estimating effects of carsharing on travel behavior
B1	**Non-overlapping support:** Non-overlapping observed characteristics of program participants and non-participants. (+)	Around 2% of the sample of employed people in the San Francisco (SF) Bay Area do not have a driver's license. These non-members do not have any equivalent members (based on the 2011-12 CHTS survey).
B2	**Differences within support**: Differences in distribution of observed characteristics within the common support. (+)	On average, members are more educated: 87% have an undergraduate degree or higher, compared to 63% of non-members. Similarly, members live predominantly in urban areas, while non-members are predominantly suburban dwellers (based on the 2011-12 CHTS survey).
B3	Selection on unobservables or **omitted variable bias**: Unobserved variables influencing both the decision to sign up for treatment and the behavioral outcome. (+++)	People with a general dislike for driving are likely to own fewer vehicles in the household and unlikely to choose carsharing over transit for their mobility requirements. However, most national and regional travel surveys collect little or no information to gauge attitudes.
B4	**Reverse causality bias or simultaneity bias**: Program participation may be caused by the variable *labeled as* the outcome, suggesting a yèd causal pathway.	Vehicle holdings (y) influence the propensity to enroll in carsharing (d). Carsharing (d), on the other hand, incentivizes reduction in vehicle holdings (y) (our initial hypothesis).
B5	**Bad conditioning**: Conditioning on a variable that represents an outcome of both the program participation and outcome variable(s) being studied (++)	Commute mode choice is known to depend upon vehicle holdings, and is likely to depend on carsharing participation. Hence, a regression of vehicle holdings on carsharing and various socio-demographic variables with commute mode choice as a covariate will bias the coefficient estimates.

Notes: (+) The examples in this table are based on differences between members and non-members on a single covariate. Similar logic applies in the k-dimensional covariate space where k represents the number of covariates, their interactions, and second-order terms as well as other transformations such as log and exponential. (++) Elwert, Winship (2014) refer to this bias as an endogenous selection bias, while referring to B1-B3 as examples of "confounding variable bias". (+++) Most papers in the literature refer to omitted variable bias (B3) when discussing self-selection bias. See for example Mokhtarian, Cao (2008).

Self-selection bias in the context of carsharing may be framed in the following way: people who self-select to become members may differ from non-members in terms of socio-economic and demographic characteristics, residential location choices, and attitudes. Our analysis of the CHTS highlighted large differences in observed characteristics between members and non-members. Members are more likely to reside in dense urban neighborhoods where carsharing is available, unlike non-members, who predominantly reside in suburban regions.

Focusing only on employed respondents in the Bay Area, we find that members are more likely to be male (61% versus 51%), highly educated (87% with a college degree or higher versus 64%), and reside in multi-unit housing complexes rather than single family homes (47% versus 17%). The potential for omitted variable bias arises because information on important confounding variables may not be collected as part of a survey. For example, CHTS measure a wide range of demographic characteristics but collect little information on attitudes towards such issues as the environment, driving, transit use, and urban living. However, such attitudes have been shown to considerably improve the ability to explain people's travel and residential location decisions (Handy, Cao and Mokhtarian 2005).

The potential for simultaneity bias in the context of the effect of carsharing on vehicle ownership based on cross-sectional data may be framed in the following way: carsharing members indeed have the incentive to reduce their vehicle holdings given the large costs of ownership in urban neighborhoods (e.g., parking), but simultaneously, people without adequate vehicle access join carsharing to enhance their mobility options (vehicle ownership influences carsharing adoption). Using a carefully designed survey administered to 9000+ members across the U.S. and Canada in 2008, Martin, Shaheen, Lidicker (2010) highlighted the complex two-way relationship between carsharing and vehicle holdings. The study asked the respondents about the circumstances under which they joined carsharing. Around 30% indicated that they joined carsharing to avoid purchasing an additional vehicle or to dispose of their vehicle, highlighting the impact of carsharing on vehicle ownership. Even more, however (around 50%) indicated that they did not own a vehicle and joined carsharing to get access to vehicles - highlighting the strength of the relation in the reverse direction.

In our carsharing research (Parts I and II), we necessarily assumed that the nature of self-selection bias arising from omitted variables (B3) – both magnitude and direction – is the same across both program participants and non-participants. In our study to estimate the effect of telecommuting adoption on travel behavior (Part-IV), we relax this assumption and explicitly measure this bias using a method called the Heckman Treatment Effect Model.

In addition to the biases B1 to B4 listed above, we researchers may induce bias by bad conditioning (B5), i.e. by controlling for a variable that represents an effect of both the program participation and outcome variables being studied (Elwert and Winship 2014; Angrist and Pischke 2008 pp. 47-49). By "controlling", we mean for example incorrectly including the variable as an explanatory variable in a parametric (example: regression) analysis. Or we may incorrectly restrict our analysis to a subsample formed on the basis of the variable (for example, we may only consider a sample of urban dwelling respondents). In our analysis, we judiciously exclude variables influenced by the outcomes being studied from both the matching process and subsequent parametric analysis to avoid bad conditioning (B5). For example, one-way commute distance is potentially an outcome of vehicle holdings – the outcome of interest in our study of effect of carsharing enrollment. As a result, we exclude the variable from any of our parametric or non-parametric methods to avoid inducing B5.

3. Methods to make causal inferences from cross-sectional data

There are a number of methods in statistical and econometric literature to address the above biases and draw causal inferences from cross-sectional observed data. Broadly, these methods may be divided into two groups. The first set of methods estimates causal effects under the assumption of unconfoundedness. That is, observed differences in characteristics, including presence and level of the program being investigated, account for the differences in the outcome of interest and selection bias is restricted to types B1 and B2 above. A second set of methods estimates causal effects when this assumption of unconfoundedness is relaxed and unobserved differences in characteristics and tastes are expected to influence both the propensity to enroll in the program as well as the outcome of interest (i.e. selection bias also arises from B3). While most of the methods seek to address selection biases B1, B2 and B3, the procedure of Structural

Equation Modeling (SEM) provides an opportunity to address the simultaneous presence of both selection bias and reverse causality bias (B4).

In this section, we provide a conceptual and accessible summary of five key methods used in the causal inference literature to make causal inferences from cross sectional data. The reader is encouraged to refer to the subsequent sections of the dissertation and other cited literature for detailed discussion.

Matching: Matching methods include either Propensity Score (PS) based matching or Mahalanobis Distance matching. In these non-parametric methods, biases B1 and B2 are addressed by matching each program participant to a non-participant with similar observed characteristics. Further, the sample is trimmed and observations without corresponding matches are removed. The causal effect is estimated by comparing the average outcomes (Y) of the participant and non-participant groups. In PS-based matching, similar characteristics include not only the observed covariates, but also their interactions and transformations like logarithmic and polynomial. The reader may refer to the following for more detailed discussions of using matching methods to draw causal inferences: Imbens, Wooldridge (2009); Ho et al. (2007); Rubin (2001).

Regression based methods: Depending upon the outcome of interest, regression encompasses two broad categories: linear regression and non-linear regressions like multinomial or ordered logistic or probit regression. In these parametric methods, biases B1 and B2 are controlled by including (also known as conditioning on) various covariates in the regression model, while estimating the effect of program participation on the outcome of interest (Y). The reader may refer to the following for more detailed discussions of using regression methods to draw causal inferences: Angrist, Pischke (2008, Chapter 3); Imbens, Wooldridge (2009).

Two-stage predictor substitution (2SPS): In the regression model estimating the outcome of interest in 2SPS, the program participation indicator is replaced by a variable that excludes variation induced by the omitted variables responsible for self-selection bias (“bad variation”). In other words, the substitute only encompasses the “good variation” that excludes omitted variable bias B3. Self-selection biases due to observed differences in characteristics (B1 and B2) are addressed by including these covariates in the regression model. The bifurcation of variation in the program participation indicator into good and bad variation is achieved using instrumental variables (IV) – variables that influence program participation but are uncorrelated with the outcome of interest. Relevant citations include Adams, Almeida, Ferreira (2009); Cerulli (2015).

We should note that the 2SPS method is an adaptation of the more commonly used two-stage least squares (2SLS) method. The adaptation is necessary because the treatment indicator is a dichotomous 0/1 indicator and not a continuous variable as is assumed by the 2SLS method.

Control Function: This method also relies on IVs to bifurcate the variation in program participation indicator into good and bad variation. However, unlike the 2SPS method, the program participation indicator is left in the regression model. Instead, a function that summarizes the bad variation (also known as the self-selection bias due to omitted variables) is explicitly included in the regression model. Unlike the 2SPS method, the CF method allows for estimation of the self-selection bias. Relevant citations include Maddala (1986); Wooldridge (2015); Mokhtarian, van Herick (2016).

Structural Equation Model: In the SEM framework, the regressions describing program enrollment and the effect of the program on the outcome of interest are simultaneously estimated. Unlike 2SPS and CF, the framework allows for feedback loops between program enrollment indicator and outcome (referred to as a non-recursive model), to estimate the influences in both directions and control for reverse causality bias (B4). Further, the errors of the two equations are allowed to co-vary, capturing the aggregated confounding influence of all omitted variables and thus controlling for omitted variable bias (B3).

In this dissertation, we discuss these methods in greater detail and highlight their strengths and weaknesses. In Part II, we combine PS-based matching with regression methods to address biases B1 and B2. In Part III, we estimate an SEM model on a matched sample (PS-based matching) to address biases B1 through B4 and compare our results with an SEM model on the full sample. In Part IV, we estimate effects and compare the results of a 2SPS and CF model to control for biases B1, B2 and B3.

4. Results

In Parts II and III of the dissertation, we investigate the effect of enrollment in carsharing on vehicle holdings. Figure 1, reproduced from Part III, highlights the role of various biases in explaining the observed differences in vehicle holdings between carsharing members and non-members. The first column gives the difference in average household vehicle holdings between non-members and members. In our sample of 8,512 San Francisco Bay Area employed respondents, the average household vehicle holdings are 1.1 and 2.0 for members and non-members respectively (hence a difference of 0.9 vehicle units).

Figure 1. Disaggregation of observed differences in vehicle holdings into causal effects and various biases

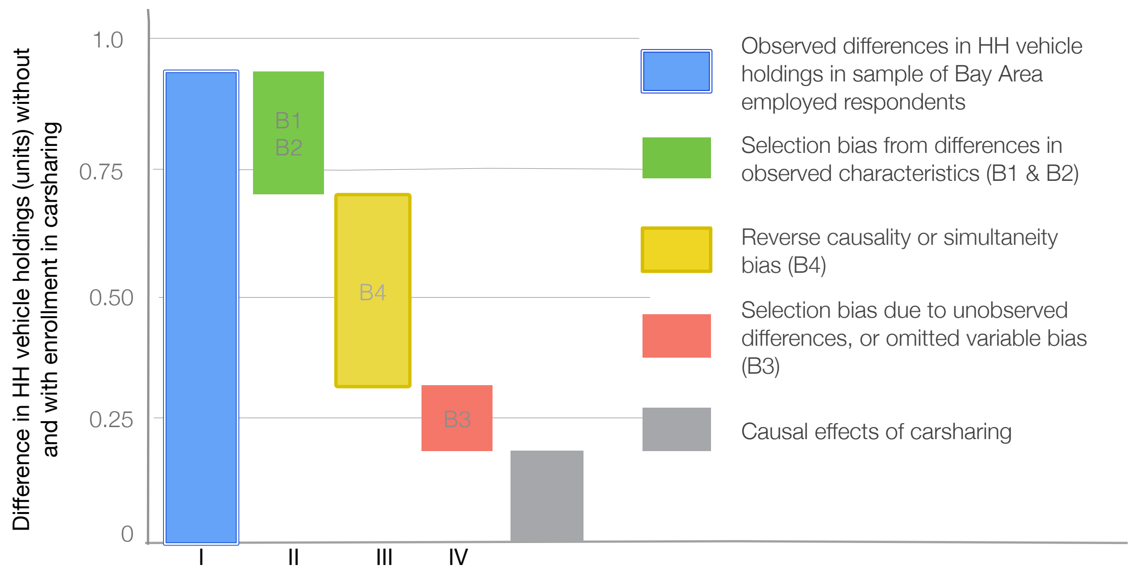

Note: The observed difference in vehicle holdings (Column I) is based on observed data in the entire sample. Columns II, III and IV are estimates based on our models.

We estimate that a difference of ~0.2 units may be attributed to observed differences in characteristics - in our study, these characteristics may be summarized as individual and household demographics, and residential and workplace neighborhood characteristics (B1 & B2). The third bar in Figure 1 (~0.4 vehicle units) may be attributed to reverse causality bias (B4) while the fourth bar (~0.15 units) may be attributed to omitted variable bias. The fifth bar, representing approximately 0.17 units, reflects the causal effect of carsharing on vehicle holdings. It is the equivalent of shedding one vehicle by about one out of every six households whose member(s) are enrolled in carsharing.

In Part IV, we look at the effects of the adoption of telecommuting on travel volumes, using the 2009-2013 waves of the UK National Travel Survey. Comparison of the observed travel volumes (Figure 2 below) indicates that telecommuters have higher mean and median annual travel volumes across all travel purposes.

Figure 2. Observed annual travel disaggregated by trip purpose and travel mode

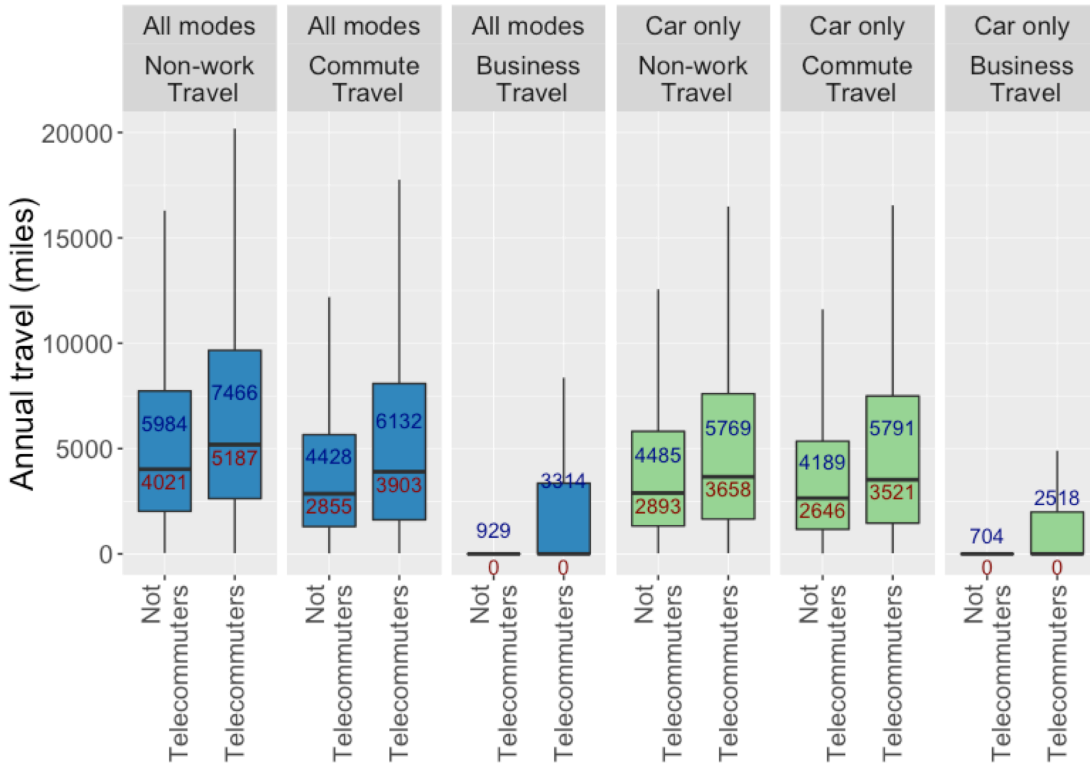

Note: The top number in blue font gives the mean, the bottom number in dark red gives the median. Sample size: 11,471 non-telecommuters and 1,475 telecommuters.

In Figure 3, we show the estimated effects of telecommuting in terms of percentage change in annual travel, disaggregated by trip purpose, mode, and current telecommuting status (to identify differences in the treatment effect on the treated and on the controls). We present the estimates based on two procedures - regression and CF. The regression method controls for biases B1 and B2 but fails to address omitted variable bias or B3. Both 2SPS and CF address B3 in addition to B1 and B2; however, the 2SPS method restrictively assumes that the nature and magnitude of selection bias arising due to omitted variables is similar for both telecommuters and non-telecommuters. The CF method relaxes this assumption and we consider it our "best" model. Nevertheless, the results of the 2SPS method are quite similar to those of the CF method; this may be partly explained by only small differences in the magnitude of the self-selection bias between the telecommuting and non-telecommuting groups (as estimated by the CF method).

Figure 3. Estimated effect of telecommuting adoption on annual travel using two methodologies – disaggregated by trip purpose, travel modes, and current telecommuting status.

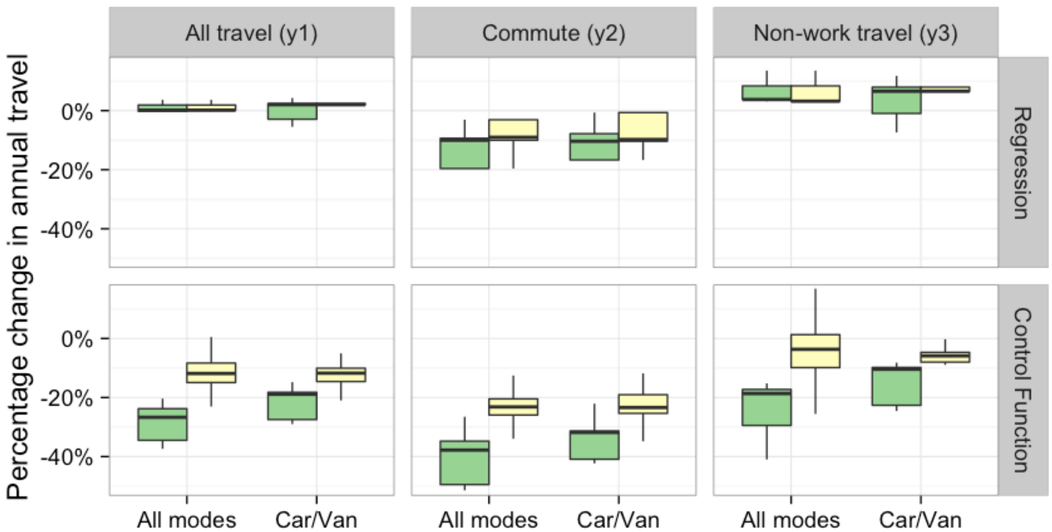

Our analysis indicates that after accounting for observed differences (the "regression" results, biases B1 & B2), the higher travel observed for telecommuters in Figure 2 fades to being (nearly-) insignificantly different from non-telecommuters. After accounting for unobserved self-selection into telecommuting (the CF results, bias B3), we find that for telecommuters, telecommuting (a) doubles the estimated reduction of commute travel due to telecommuting (to ~23%); (b) switches the estimated effect on non-work travel from positive to negative (but insignificant in both cases); and (c) changes the estimated effect on total travel (except business) from 0 to 12% reduction.

5. Conclusions

In this dissertation, we develop a methodological framework that capitalizes on existing procedures in the statistical and econometric literature to draw causal inferences from cross-sectional observed data. Our analysis highlights the presence of, and hence need to control, four broad biases – self-selection bias due to observed differences in characteristics (B1 & B2 of Table 1), self-selection bias due to omitted variables (B3) and reverse causality bias (B4). In carsharing, we find that 80% of the observed differences in average household vehicle holdings between carsharing non-members and members may be explained by the biases listed above. In telecommuting, we find that the higher observed travel volumes of telecommuters compared to non-telecommuters are explained by observed differences in characteristics (B1 and B2). Further, controlling for omitted variables, leads us to conclude that telecommuting leads to substantial reduction in travel volumes.

Our results for both carsharing and telecommuting are consistent with earlier studies, although nearly all those, unlike our study, are based on purpose-built proprietary surveys with questionnaires that rely on retrospective reporting (hence panel studies) and are explicitly designed to evaluate effects of either of the two programs. With respect to the panel studies on telecommuting, Kim, Choo, Mokhtarian (2015) and Andreev, Salomon, Pliskin (2010) conclude that most of them find a substitution effect, i.e. telecommuting leads to less travel, and the results are usually substantial and statistically significant. These conclusions hold for both commute as well as non-commute travel. In carsharing, Martin, Shaheen, Lidicker (2010) and Firnkorn, Müller (2012) estimated reductions of 0.18-0.26 and 0.05-0.11 vehicles per household respectively.

We hope that our research will encourage more studies to rigorously infer causal effects from publicly available national and regional travel surveys. Nevertheless, researchers should be cognizant of the limitations of the methods, and the strong assumptions necessary for drawing causal inferences from cross-sectional observed data with potential for presence of multiple biases. We elaborate on these assumptions and limitations in all three parts of the dissertation. Here we touch upon one key limitation – the necessity for what is known as an instrumental variable, or IV.

Both the Two-stage Predictor Substitution (2SPS) and Control Function (CF) methods require the availability of IV(s). A Structural Equation Model that is both non-recursive (to control for reverse causality or B4) and having correlated errors (to control for omitted variables or B3) also requires IV(s) for estimation. IVs are exogenous variables that fulfill two key criteria.

First, they should "sufficiently" explain the variation in the indicator capturing program participation – in other words, they should influence the propensity to enroll in the program. Second, the sole channel through which the IV(s) affect the outcome variable *y* of interest should be through the program participation indicator. In other words, the IV should not affect outcome directly even though the empirical correlation between the IV and *y* is unlikely to be zero because of the likely correlation of both *y* and the IV with the program participation indicator.

Unfortunately, it is difficult to obtain proper instruments that satisfy the above two conditions – especially in general purpose observational data like national and regional travel surveys. In Part III of the dissertation, we depended upon an IV that was synthetically estimated by us based on the sample distribution and not originally collected as part of the CHTS. Although we adopted very strict criteria to estimate the indicator, the synthetic nature of the variable gives us reason for caution in adopting our estimates of causal effects. In Part IV of the dissertation, we investigate the appropriateness of an IV used earlier in the literature to capture the propensity to adopt telecommuting. Zhu (2013) and Zhu, Mason (2014) instrumented the telecommuting indicator with a variable indicating frequency of use of internet at home in their study of the effect of telecommuting on one-way commute distance and frequency of work and non-work trips. However, it is unclear whether this indicator can effectively discriminate between telecommuters and non-telecommuters, given the near ubiquity of internet access including broadband internet access. For example, in our dataset of employed non-home-based workers from the UK NTS, nearly 93% of respondents indicated availability of broadband at home. Among telecommuters, the ratio was 99%. More importantly, we believe that the indicator may not meet the second criterion of a good IV. Mokhtarian (2009) identifies various ways in which internet availability and usage can facilitate both increases and decreases in travel, and not all these uses are mediated by telecommuting. For example, internet access provides an opportunity to substitute physical shopping with online shopping, thereby eliminating the need for some shopping trips. Given the existence of direct pathways through which internet access affects non-work travel, as well as the near ubiquity of access, we do not consider this variable a suitable IV.

In the dissertation, we have identified a number of other limitations that may be specific to individual methods used in this study. The discussion of these limitations highlights the difficulties in making causal inferences regarding programs with bi-directional effects, using cross-sectional observational data. However, randomized experiments or collection of long-term longitudinal data are quite rare due to cost and feasibility considerations. In their absence, the methods used in this study may offer a ballpark estimate of causal effects, provided analysts and readers are well aware of the limitations.

6. References

Adams, R., Almeida, H., Ferreira, D.: Understanding the relationship between founder–CEOs and firm performance. Journal of empirical Finance **16**(1), 136-150 (2009)

Andreev, P., Salomon, I., Pliskin, N.: Review: State of teleactivities. Transportation Research Part C: Emerging Technologies **18**(1), 3-20 (2010)

Angrist, J., Pischke, J.-S.: Mostly harmless econometrics: An empiricist's companion. Princeton University Press, (2008)

Cerulli, G.: Econometric evaluation of socio-economic programs: Theory and applications, vol. 49. Springer-Verlag Berlin Heidelberg, (2015)

Elwert, F., Winship, C.: Endogenous Selection Bias: The problem of conditioning on a collider variable. Annual Review of Sociology **40**, 31-53 (2014). doi:10.1146/annurev-soc-071913-043455

Firnkorn, J., Müller, M.: Selling mobility instead of cars: new business strategies of automakers and the impact on private vehicle holding. Business Strategy and the Environment **21**(4), 264-280 (2012). doi:10.1002/bse.738

Handy, S., Cao, X., Mokhtarian, P.: Correlation or causality between the built environment and travel behavior? Evidence from Northern California. Transportation Research Part D: Transport and Environment **10**(6), 427-444 (2005). doi:10.1016/j.trd.2005.05.002

Heckman, J.J., Ichimura, H., Smith, J., Todd, P.: Sources of Selection Bias in evaluating social programs: An interpretation of conventional measures and evidence on the effectiveness of Matching as a program evaluation metho. Proceedings of the National Academy of Sciences **93**(23), 13416-13420 (1996)

Ho, D.E., Imai, K., King, G., Stuart, E.A.: Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. Political Analysis **15**(3), 199-236 (2007)

Imbens, G.W., Wooldridge, J.M.: Recent Developments in the Econometrics of Program Evaluation. Journal of Economic Literature **47**(1), 5-86 (2009)

Kim, S.-N., Choo, S., Mokhtarian, P.L.: Home-based telecommuting and intra-household interactions in work and non-work travel: A seemingly unrelated censored regression approach. Transportation Research Part A: Policy and Practice **80**, 197-214 (2015)

Maddala, G.: Limited-dependent and qualitative variables in econometrics. vol. 3. Cambridge University Press, (1986)

Martin, E., Shaheen, S.A., Lidicker, J.: Impact of carsharing on household vehicle holdings. Transportation Research Record: Journal of the Transportation Research Board **2143**(1), 150-158 (2010). doi:10.3141/2143-19

Mokhtarian, P.: If telecommunication is such a good substitute for travel, why does congestion continue to get worse? Transportation Letters **1**(1), 1-17 (2009)

Mokhtarian, P.L., Cao, X.: Examining the Impacts of Residential Self-selection on Travel Behavior: A Focus on Methodologies. Transportation Research Part B: Methodological **42**(3), 204-228 (2008). doi:DOI: 10.1016/j.trb.2007.07.006

Mokhtarian, P.L., van Herick, D.: Quantifying residential Self-Selection effects: A review of methods and findings from applications of Propensity Score and Sample Selection approaches. Journal of Transport and Land Use **9**(1) (2016)

Rubin, D.B.: Using propensity scores to help design observational studies: application to the tobacco litigation. Health Services and Outcomes Research Methodology **2**(3-4), 169-188 (2001). doi:DOI:10.1023/A:1020363010465

Rubin, D.B.: Comment: The design and analysis of gold standard randomized experiments. Journal of the American Statistical Association **103**(484), 1350-1356 (2008)

Wooldridge, J.: Control function methods in applied econometrics. Journal of Human Resources **50**(2), 420-445 (2015)

Zhu, P.: Telecommuting, household commute and location choice. Urban Studies **50**(12), 2441-2459 (2013)

Zhu, P., Mason, S.G.: The impact of telecommuting on personal vehicle usage and environmental sustainability. International Journal of Environmental Science and Technology **11**(8), 2185-2200 (2014)

PART-II

Paper 1: The effect of carsharing on vehicle holdings and travel behavior: a Propensity Score and Causal Mediation analysis of the San Francisco Bay Area

Published in Research in Transportation Economics 52 (2015): 46-55. DOI: 10.1016/j.retrec.2015.10.010

THE EFFECT OF CARSHARING ON VEHICLE HOLDINGS AND TRAVEL BEHAVIOR: A PROPENSITY SCORE AND CAUSAL MEDIATION ANALYSIS OF THE SAN FRANCISCO BAY AREA

Gouri Shankar Mishra (Corresponding author)
Institute of Transportation Studies, University of California, Davis
gsmishra@ucdavis.edu

Regina R. Clewlow
Precourt Energy Efficiency Center, Stanford University
rclewlow@stanford.edu

Patricia L. Mokhtarian
School of Civil and Environmental Engineering, Georgia Institute of Technology
patmokh@gatech.edu

Keith F. Widaman
Graduate School of Education, University of California, Riverside
keith.widaman@ucr.edu

Abstract

We examine the impacts of carsharing on travel behavior utilizing a San Francisco Area subsample of the 2010-2012 California Household Travel Survey. We control for self-selection bias due to differences in observed characteristics of the respondents using propensity-score based matching. We find that vehicle holdings of carsharing members are substantially and significantly lower than for non-members with similar characteristics in terms of individual and household demographics and built environment features of both residential and job location. These differences increase as the propensity to enroll in carsharing programs increases. A latent construct, which measures the propensity to own or utility from owning vehicles and rises with numbers of vehicles owned, is lower for members by 0.3 to 1.3 standard deviations relative to non-members. Members are also likely to walk, bike, and use transit more frequently than non-members. However, these differences are relatively minor and tend to be statistically non-significant. Future research should control for self-selection bias arising from differences in unobserved characteristics of respondents, as well as simultaneity bias whereby decisions concerning vehicle ownership both influence and are influenced by the decision to join carsharing programs.

Keywords: Self-selection bias, carsharing, shared use mobility, propensity score matching, vehicle ownership, alternative modes

1. Introduction

The significant growth in shared use mobility alternatives (e.g., carsharing, on-demand ride services) has prompted policymakers and the automotive industry to consider how these services impact travel behavior, vehicle ownership, and associated energy and greenhouse gas (GHG) emissions of the transportation sector. In recent decades, the availability of carsharing services such as Zipcar™ has spread throughout many cities, providing as-needed access to vehicles, typically on an hourly basis. Previous carsharing research, based largely on member surveys, suggests that carsharing leads to vehicle shedding, reduced vehicle miles traveled (VMT), and increased use of public transit, walking, and biking (Cervero, Golub and Nee 2007; Martin and Shaheen 2011a; Sioui, Morency and Trépanier 2013; Shaheen, Cohen and Chung 2009), although studies differ in terms of the magnitude of impact (Tal 2009).

One of the key limitations of previous carsharing work is that the adoption of carsharing is likely coupled with the decision to live in a dense, urban area, which in itself is known to have a significant impact on travel behavior (Handy, Cao and Mokhtarian 2005; Bhat and Eluru 2009). In addition, other factors that pre-dispose an individual to adopt carsharing may also shape her travel decisions, the primary one being limited access to household vehicles. Failure to account for these prior propensities may lead to overestimating the effect that carsharing would have on travel behavior if adopted by someone without those propensities, i.e. if policies promoting carsharing led to its adoption across a broader segment of society. This study presents the first attempt to account for self-selection bias in carsharing – specifically, to account for the potential pre-disposition towards carsharing enrollment in order to identify more accurately the effect of carsharing on travel behavior.

This study utilizes data from the 2010-2012 California Household Travel Survey (CHTS). The CHTS collects detailed information about household demographics and travel activity for the purposes of modeling statewide and regional travel and GHG emissions (Caltrans 2013). Sampling and weighting methods are employed to match statewide household population distributions on key demographic variables. In the 2010-2012 survey, respondents were asked whether they were members of a carsharing organization. Around 800 of the 84,000 individuals 18 years or older indicated they were members, and 80% of these carshare members were employed.

We adopt a non-parametric matched-sampling procedure to identify a control group that is statistically balanced on various observed socio-economic traits and residential location choices, and whose travel behavior may be compared with that of carsharing members (treatment group). Matched sampling is well established in diverse disciplines such as political and legal studies (Ho and Rubin 2011), epidemiology and medical research (Weitzen et al. 2004), and economics (Caliendo and Kopeinig 2008), as a statistical method to make causal and counterfactual inferences from observational data. The goal of matched sampling (or simply matching) is to balance the distribution of observed confounding covariates between the control and treated groups so that the resulting differences in outcomes between the groups may be attributed to the treatment under study. Further, matching reduces model dependency in parametric analysis and improves the precision of estimated coefficients (Ho et al. 2007).

The remainder of this paper is structured as follows. In Section 2, we review key studies on carsharing. We discuss methods in Section 3 and results in Section 4. In Section 5, we discuss the limitations of this analysis, future work, and potential policy implications of this research.

2. Background

2.1. Carsharing

A number of studies have sought to examine the causal impact of carsharing membership on travel behavior. The first major such study involved multiple surveys, between 2001 and 2005, of City CarShare members, a non-profit carsharing organization based in the Bay Area (Cervero and

Tsai 2004; Cervero, Golub and Nee 2007). The control group for the study consisted of respondents who demonstrated interest in joining City CarShare, but who did not join as members for various reasons, including service unavailability in their neighborhoods. Adopting a difference-in-difference (DID) of means evaluation method for summary statistics, the study found that between 2001 (a few weeks before membership or treatment) and 2005, carsharing membership led to a decline of weekday total travel and vehicle miles traveled (VMT) by 4.50 and 4.26 miles per day respectively. Both members and non-members reduced their VMT and fuel consumption over the 5-year period – potentially due to rising oil prices – but the reduction was steeper for members.

One of the largest studies involved a survey of 9000+ members across the U.S. and Canada in 2008 to assess the impact of carsharing (Martin, Shaheen and Lidicker 2010; Martin and Shaheen 2011a). Members reported their travel behavior both currently and retrospectively prior to joining carsharing. The study found that at an aggregate level, members increased their use of public transit after joining carsharing; however there was significant variability in the results. An increase in walking, biking, and carpooling, compared to levels prior to enrollment in carsharing, was more definitive. Further, carsharing members reduced their household vehicle holdings – from an average of 0.47 to 0.24 vehicles. In the absence of comparison with a control group and statistical control of various confounding variables, including potential relocation to an urban area prior to joining carsharing (together with the lower reliability of retrospective reporting), it is not possible to determine whether the entirety of these observed behavioral shifts can be attributed to carsharing membership.

According to the research by Martin, Shaheen, Lidicker (2010), demand for carsharing comes largely from households with limited access to cars, and carsharing constitutes the primary access to a car for most active members. Martin, Shaheen (2011b) found that most members they surveyed (62%) did not own a car at the time of enrollment in carsharing. To put this number in perspective, only 10% of households residing in urban regions, which includes cities and their suburbs, in the U.S. do not own a car (U.S. Census Bureau 2014). Analyzing the household travel surveys in Montreal, Klincevicius, Morency, Trépanier (2014) similarly found that household car ownership was negatively correlated with carsharing service availability in a neighborhood (number of shared cars in a 500-m radius) after controlling for various socio-economic confounders. Stillwater, Mokhtarian, Shaheen (2009) found a negative correlation between demand for carsharing at a pod location and average car ownership levels in the neighborhood.

A more recent paper by Sioui, Morency, Trépanier (2013) compared members of Montreal-based Communauto with the larger Montreal population. The study was facilitated by a purpose-built survey of members and an independent large-scale household travel survey conducted around the same time. The study found that members are less likely to own vehicles – around 90% of member households did not own vehicles compared to 34% of the general population. In general, the share of non-motorized trips was higher among carsharing households than among the rest of the population. Comparing only those without vehicles, however, carsharing households drove more and used public transit less than non-member households, suggesting either that carsharing households differed in unobserved ways from their non-member counterparts, or that carsharing enabled them to shift some travel from public transit to automobile.

Whereas the above studies were all based on observed (or revealed) behavior, at least two studies assessed likely changes in behavior if respondents enrolled in a yet-to-be introduced carsharing service. Zhou, Kockelman (2011) and Firnkorn, Müller (2011) found that respondents would likely reduce their vehicle holdings if they enrolled in car2go in Austin, Texas and Ulm, Germany respectively.

This paper builds on previous research to assess the impact of carsharing on travel behavior in two specific ways. It uses a well-defined statistical method to control for selection bias (discussed in detail in the next section), and it undertakes a causal mediation analysis to decompose the pathways through which carsharing likely affects short-term travel behavior such as the frequency of transit, walk, and bike trips, as well as daily driving distance.

2.2. Self-selection bias

Borrowing notation from the Rubin Causal Model framework, the difference in average travel outcomes (Y) between members (CS=1) and non-members (CS=0) can be decomposed into a causal effect and a self-selection bias

$$\Delta\mu = E\left[Y_{1,i}|CS = 1\right] - E\left[Y_{0,i}|CS = 0\right] \tag{1A}$$

$$= \left(E\left[Y_{1,i}|CS = 1\right] - E\left[Y_{0,i}|CS = 1\right]\right) + \left(E\left[Y_{0,i}|CS = 1\right] - E\left[Y_{0,i}|CS = 0\right]\right) \tag{1B}$$

where the *i* subscripts individual members, the expectation is taken over *i,* and $\Delta\mu$ is the observed difference in average outcomes $E\left[Y_{1,i}\right] - E\left[Y_{0,i}\right]$. The first two terms on the right-hand side of Eq. (1B) give the average causal effect of carsharing on members (CS=1), also called the average treatment effect on the treated (ATT). It is the expected difference between the observed outcome ($Y_{1,i}$) and a counterfactual outcome ($Y_{0,i}$) had the member not enrolled in carsharing. The last two terms give the self-selection bias, as the expected difference between the counterfactual outcome of members had they not enrolled ($Y_{0,i}|CS = 1$), and the observed outcome of non-members ($Y_{0,i}|CS = 0$).

Self-selection bias in the context of carsharing may be framed in the following way: people who self-select to become members may differ from non-members in terms of socio-economic and demographic characteristics, residential location choices, and attitudes. For example, compared to non-members, members are less likely to own a car even before membership, are younger and more educated, and live in smaller households in urban centers (Martin, Shaheen and Lidicker 2010; Sioui, Morency and Trépanier 2013; Loose 2010).

Parametric estimates (linear regression or discrete choice models) of causal effects from cross-sectional observed studies may be undertaken using the entire sample. In these models, membership may be represented as a binary indicator, and included as an explanatory variable along with other covariates to explain the travel behavior outcome under study. The coefficient of the carsharing indicator may be construed as the causal effect on the outcome under study. However, if there are large differences in the distributions of the various explanatory variables (confounders or covariates) of the outcome under study between members and non-members (control group), then self-selection bias is present (Heckman et al. 1996), and the estimated coefficient will be biased. Specifically, Imbens, Rubin (2015 p 277) indicate that if the mean of program participants and non-participants for "important" covariates are more than one-quarter or at most one-half standard deviations apart, then the regression estimates will be seriously biased. In addition to bias, differences in the distribution of covariates may lead to increased "model dependency", meaning that the estimated causal impacts presented are not very robust with respect to variations in model specification (Ho et al. 2007).

Our initial analysis of the CHTS highlighted large differences in observed characteristics or covariates between members and non-members as summarized in Table 3 and Supporting Information (SI) Section S2. Members are more likely to reside in dense urban neighborhoods where carsharing is available, unlike non-members, who predominantly reside in suburban regions. Focusing only on employed respondents in the Bay Area, we find that members are more likely to be male (62% versus 52%), highly educated (87% with a college degree or higher versus 65%), reside in multi-unit housing complexes rather than single family homes (47% versus 15%), and less likely to commute to work by car (32% versus 82%). These findings are largely consistent with observations in the carsharing literature.

Table 1. Key variables used in this study

Variable	Description
Outcomes [+]	
# of HH vehicles (Y_A & M_A)	Number of vehicles owned / leased by household (HHVeh). Modeled as a moderator to measure impact on other outcomes listed below.
Vehicles per driver (Y_B & M_B)	An alternate indicator of vehicle holdings. Equal to # of household vehicles divided by # of household members with driver's license (VehDriver).
Transit trips per week (Y_C)	Number of one-way transit trips made by respondent in the week prior to survey date.
Non-motorized trips per week (Y_D)	Number of times a person walked or biked outside including trips taken for exercise in in the week prior to survey date.
Treatment Indicator	
CS	Indicator of carsharing membership (treatment).
Confounders (Household, or HH, level)	
Child	Number of children (16 years or younger) in HH
HH employees	Number of HH members reporting to be workers
Income	Household (HH) income: 5 categories ranging from \$50,000 or less to \$250,000 or more
Own	Home owned (1) or rented (0)
Residence type	Residence type: 3 categories (single family, multi-unit housing, others including mobile home and boats)
Residential neighborhood accessibility Work neighborhood accessibility	Composite indices computed by www.walkscore.com™ reflecting the accessibility of five categories of destinations for the residential / workplace census tracts: educational (e.g., schools), retail (e.g., grocery, drug, and convenience), food (e.g., restaurants), recreational (e.g., parks) and entertainment (e.g., theaters). Ranges from 0 to 100 (high accessibility). See SI Section 1
Residential neighborhood density Work neighborhood density	Composite indices reflecting residential and population density of the census tracts in which subject resides / works respectively. The indicator was estimated using exploratory factor analysis of five manifest variables from U.S. Census American Community Survey (ACS) 5-year estimates (See SI Section 1). Ranges from -2 (low density) to +4 (high density)
Confounders (Individual level) [++]	
Age	Age in years
Commute distance	Commute distance in miles. Gives the driving distance from centroid of residential census tract to that of work location tract and calculated using MapQuest. If same census tract, we assume a distance of 2.5 miles. The distribution is lognormal; hence we consider the natural logarithm
Disability	Indicator of temporary or permanent physical condition or disability that makes travel outside home difficult
Education	Education level reflecting highest degree or level completed. 4 categories from "High school or less" to "Graduate degree"
Female	Gender: male (0) or female (1)
License	Indicator for a valid driver's license (DL)
Occupation	Occupation, aggregated to 3 categories from around 25 in CHTS: white-collar (CHTS "occup" code 11-31), blue-collar (32-53), and military & others (55 & 97)
# days commuted	Days in a week person commutes to work location (1 to 7)

Notes: (+) Modeled as ordinal variables. (++) Information on respondent's primary commute mode is available in CHTS but we excluded it from both propensity score model for matched sampling as well as subsequent parametric analysis on the matched sample. The issue is discussed further in Section 3.2.

2.3. Matched sampling

We adopt a matched sampling approach to address the self-selection bias arising due to differences in observed covariates and to estimate the effects of carsharing. To control for self-selection using a matched sampling method, one or more control units are selected for each treatment unit such that the units have the same or similar values for various observed covariates. In other words, through targeted trimming, the distribution of covariates in the control group is reshaped to closely match the distribution of covariates in the treatment group, thereby minimizing any selection bias arising due to observed covariates and making the outcome variables independent of carsharing membership conditional on **X** i.e. (Y_{1i}, $Y_{0,i}$ $\perp\!\!\!\perp$ CS_i | X_i).

As a result of the conditional independence assumption, $E[Y_{0,i}|X_i, CS = 1] = E[Y_{0,i}|X_i, CS = 0] = E[Y_{0,i}|X_i]$. This allows us to reduce Eq. (1B) to the following:

$$\Delta\mu = E[Y_{1,i}|X_i] - E[Y_{0,i}|X_i] = E[Y_{1,i} - Y_{0,i}|X_i] \qquad (2)$$

The distribution of covariates in the matched sample matches the distribution of the carsharing members before matching. As a result, the treatment effects estimated in Eq. (2) may be interpreted as ATT.

There are multiple approaches to matching. In this study, we match on the propensity score $e_i(X_i) = P(CS_i = 1|X_i)$, defined for each subject as the probability of joining carsharing (CS=1) given the observed **X**. Rosenbaum, Rubin (1983) show that at each value of the "true" propensity score (PS), the distribution of the covariates **X** is the same in the treated and control groups regardless of the dimension of **X**. Thus subjects in the control group with $e_i(X_i)$ close to a treatment subject will have similar values on **X** and are hence part of the matched sample (subject to other specifications).

We note at the outset that matched sampling addresses self-selection bias arising due to observed covariates, but not unobserved covariates. For example, self-selection bias arising due to attitudes towards carsharing cannot be fully addressed in matched sampling because the CHTS dataset does not enable us to directly gauge the attitudes of individuals. In the context of a parametric model to explain the influence of carsharing on travel behavior, this may be referred to as an omitted variable bias because these omitted variables (attitudes) correlate with the indicator representing carsharing membership. The issue is discussed further in Section 5.

2.4. Effect on vehicle holdings

As mentioned before, we adopt two alternatives to model vehicle holdings: (a) Number of household vehicles (Y_A, HHVeh), and (b) Number of vehicles per driver in the household (Y_B, VehDriver). We aggregated household vehicle holdings (Y_A) into four ordinal levels – 0, 1, 2, and 3 or more vehicles per household. Vehicles per driver (Y_B) is conceptually a continuous variable, however empirically its kernel density (SI Section 3) reveals a multimodal distribution. As a result, we transformed the variable into a categorical ordinal variable with three levels - less than one, one, and more than one vehicle per driver in the household.

In both cases, we model a latent variable U* which may be considered as a measure of the propensity to own vehicles and rises with numbers of vehicles owned, and modeled as a function of all the respondent characteristics listed in Table 1 (Greene and Hensher 2010; Bhat and Pulugurta 1998). The function for latent propensity or utility for individual *i* may be represented as:

$$U_i^* = \beta_1^a CS_i + \boldsymbol{\gamma^a X_i} + \boldsymbol{\varepsilon_i} \qquad \textbf{(3A)}$$

where CS is the carsharing indicator; **X** is a vector of all the other covariates. The "*" reflects the latent (unobserved) nature of the utility, unlike the covariates in the model. The causal effect of carsharing on vehicle holdings is estimated by β_1^a. The utility or propensity U* may be mapped to the observed discrete levels of vehicle ownership (HHVeh)

$$Y_{A,i} = k \Rightarrow \text{HHVeh} = K_A \quad \text{if } \tau_{k-1} \leq U_i^* < \tau_k;\ 0 \leq k \leq 3;\ \tau_{-1} = -\infty\ \&\ \tau_3 = +\infty \tag{3B}$$

where τ_k reflects the thresholds assumed constant across all respondents. The ordinal levels of HHVeh are represented as K_A and detailed in Section 2.2. We estimated the above using an ordered probit model assuming a normal distribution with variance of 1 for the error term $\boldsymbol{\varepsilon_i}$ in Eq. (3A). The probability that a respondent chooses a particular level of vehicle ownership can then be represented as:

$$\Pr(Y_{A,i} = k|\boldsymbol{X_i}, CS_i) = \Phi(\tau_k - \beta_1^a CS_i + \boldsymbol{\gamma^a X_i}) - \Phi(\tau_{k-1} - \beta_1^a CS_i + \boldsymbol{\gamma^a X_i}) \tag{3C}$$

For k = 0 and hence K_A = zero HH vehicles, the probability may be expressed as $\Pr(Y_{A,i} = 0|\boldsymbol{X_i}, CS_i) = \Phi(\tau_0 - \beta_1^a CS_i + \boldsymbol{\gamma^a X_i})$. Similarly, for k = 3 and K_A=3 or more HH vehicles, the probability may be expressed as $\Pr(Y_i = 3|\boldsymbol{X_i}, CS_i) = 1 - \Phi(\tau_2 - \beta_1^a CS_i + \boldsymbol{\gamma^a X_i})$. A similar set of equations (3B & 3C) may be specified for VehDriver with the only change being that there is one less categorical level and hence one less threshold.

The difference in probability of belonging to a particular category K_A (and similarly, K_B) between carsharing members and non-members may be expressed in the following way:

$$\Delta\Pr(Y_A = k|\boldsymbol{X} = \boldsymbol{\bar{X}}) = \Pr(Y_A = k|\boldsymbol{\bar{X}}, CS = 1) - \Pr(Y_A = k|\boldsymbol{\bar{X}}, CS = 0) \tag{3D}$$

Eq. (3D) indicates that the marginal effect of carsharing membership is computed while holding other covariates at their mean levels.

2.5. Effect on frequency of transit and walk and bike trips

The travel behavior outcomes may be divided into decisions made on a longer timescale – in our case household vehicle holdings (Y_A,Y_B); and those decisions made on a shorter timescale – weekly transit trips (Y_B), and weekly non-motorized (walk and bike) trips (Y_C). We do not consider measures of vehicle usage because CHTS collects driving distances only for the day of the survey and the measure is sensitive to whether the respondent commuted to work on that day or not. Around 30% of both carsharing members and control respondents in the matched sample did not commute to work on the day of the survey.

It may be argued that carsharing affects short-term travel decisions through at least two pathways – directly, and indirectly through its longer-term impact on vehicle holdings (Y_A). For example, carsharing enhances the mobility options available to members, especially those who do not own a car, possibly encouraging more driving at the expense of transit and non-motorized trips (direct effect). On the other hand, carsharing may also influence an overall reduction in vehicle holdings, which in turn may reduce driving and increase travel by other modes (mediated effect). It is important to investigate the relative magnitudes of these two theoretically opposite effects, so as to understand the constituents of the *net* impact of carsharing on the use of transit and non-motorized modes.

As in the case of vehicles per driver, kernel density graphs of number of weekly transit and non-motorized trips suggest the lognormal distribution, however, their logarithmic transformation indicates a bi-modal distribution with the first mode close to zero (SI Section 3). As a result, we transformed them into ordinal variables with four categories each representing 0, 1-7, 8-14, 15 or more trips per week. These categories are equivalent to zero trips, one trip per day or less, 1-2 trips per day, and finally 3 or more trips per day.

Representing the mediator (Y_A or M_A) as $M_{0,i}$ and $M_{1,i}$ for individual *i* under the alternative program enrollment states (0 and 1 respectively), and noting that Y here represents either Y_C or Y_D, we can restate Eq. (2) as the following:

$$\Delta\mu = E\left[(Y_{1,i}(M_{1,i}) - Y_{0,i}(M_{0,i})\right] \tag{4}$$

The total effect can then be linearly decomposed into direct and indirect (or mediated) effects (Imai, Keele and Tingley 2010) as represented by the following set of equations:

$$\Delta\mu_I = E\left[(Y_{1,i}(M_{1,i}) - Y_{1,i}(M_{0,i})\right] \tag{5A}$$

$$\Delta\mu_D = E\left[(Y_{1,i}(M_{0,i}) - Y_{0,i}(M_{0,i})\right] \tag{5B}$$

$$\Delta\mu = \Delta\mu_I + \Delta\mu_D \tag{5C}$$

Eq. (5A) gives the mediated or indirect effect, which is the change in the outcome corresponding to a change in the mediator from the value that would be realized under the control condition, i.e., $M_{0,i}$, to the value that would be observed under the treatment condition, i.e., $M_{1,i}$, while holding the membership (treatment) status constant at CS=1. If $M_{1,i} = M_{0,i}$, then the treatment has no effect on the mediator and the mediated effect would be zero. Eq. (5B) gives the direct effects, which capture all other pathways through which carshare membership impacts the outcome under study. The term $Y_{1,i}(M_{0,i})$ is a counterfactual and not observed in the dataset.

To decompose the total effects of carsharing on outcomes Y_C and Y_D into direct and mediated effects, we adopt the multi-regression approach method detailed by MacKinnon, Fairchild, Fritz (2007) who also refer to it as the Baron and Kenny approach. The corresponding path model is depicted in Table 2. Similar to the latent utility from vehicle holdings modeled in Eq. 3A, the latent propensity to take more transit (Y_C) or non-motorized (Y_D) trips may be estimated as the following (subscript *i* not included for notational simplicity):

$$V^*_{C\,orD} = \beta_0^b + \beta_1^b CS + \boldsymbol{\gamma^b X} + \boldsymbol{\epsilon_i} \tag{6A}$$

$$W^*_{C\,orD} = \beta_0^c + \beta_1^c CS + \beta_2^c U^* + \boldsymbol{\gamma^c X} + \boldsymbol{\omega_i} \quad . \tag{6B}$$

As before, U* represents the latent propensity to own more vehicles (the mediator). V* and W* are latent propensities to take transit or non-motorized modes more often, modeled without and with inclusion of the utility resulting from the mediator (U*) respectively. Equations (6A) & (6B) are estimated separately for both transit and non-motorized trip frequency. The mapping of observed ordinal outcomes to the continuous latent utility V* or W* may be undertaken in a similar fashion as in Eq. (3B).

Table 2. Multi-regression estimation of causal impacts of carsharing membership (CS)

Step 1: Impact on longer term travel behavior	Step 2: Impact on shorter term travel behavior	
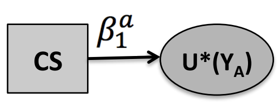	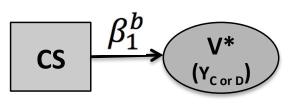	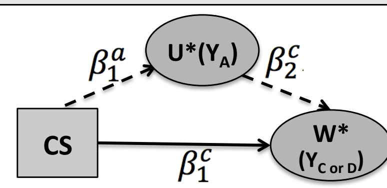
Notes: U*(Y_A) represents the unobserved or latent utility from greater vehicle holdings. β_1^a represents the impact of carsharing (CS) on U*(Y_A).	*Notes*: V* and W* represent the latent propensities to take more transit (Y_C) and non-motorized modes (Y_D), estimated without and with the presence of the mediator U*(Y_A) respectively. The dashed lines indicate the indirect effect of CS on W*, while the solid lines indicate the direct effect. β_1^b and β_1^c represent the total & direct effect, respectively, of carsharing (CS) on the utilities associated with Y_C and Y_D. The mediated effect may be estimated as ($\beta_2^c . \beta_1^a$). The confounding variables (X) are not shown for simplicity.	

The indirect or mediated effect is equal to the product $\beta_2^c . \beta_1^a$ (MacKinnon, Fairchild and Fritz 2007; Muthén 2011).

For the regressions in Eqs. 3 and 6, we estimated cluster-robust standard errors with clustering on the residential census tracts. Residents in the same neighborhood are likely to have similar travel behavior and other socio-economic characteristics. Clustering also arises because the attributes of a census tract – residential accessibility and density scores – do not vary across individuals within each tract (Angrist and Pischke 2008 Chapter 8; Moulton 1990).

3. Methods

3.1. CHTS Data

This study focuses on the San Francisco Bay Area, which contains 50% of the respondents statewide who indicated being enrolled in carsharing. We restricted our analysis to employed respondents who were 18 years or older, and who had identified a fixed primary work location outside their homes. Further, for the sake of simplicity, we dropped cases with incomplete data on the key covariates (identified in Table 1), thereby implicitly assuming missingness completely at random. The final sample, prior to matching, consisted of 8299 respondents including 241 carsharing members.

3.2. Matched Sampling

We established a matched sample based on propensity scores. Propensity scores for each individual were calculated using a logit model based on all variables identified as potential confounders. The objective of the propensity score model is, counterintuitively, not to successfully and parsimoniously predict probabilities, but to efficiently control confounding by balancing covariates. Thus, methods and diagnostic tests ordinarily applied for successful model building should be avoided during specification of the propensity score model – collinearity tests; goodness-of-fit statistics; inclusion of variables that influence treatment assignment but not outcome (e.g., availability of shared vehicles, corporate subsidies, etc.); and iterative model-building algorithms (e.g., forward stepwise regression) (Brookhart et al. 2006; Austin 2011; Weitzen et al. 2004). Rather, the focus should be to include all potential confounders; for example in Rubin (2001), 50 variables were included to calculate propensity scores to ensure that balance was achieved in all observed potential confounders. Outcome variables (Y_A to Y_D) are not included in the propensity score model (Ho et al. 2007). Further, variables that are influenced by the outcome variables of interest – specifically commute mode choice, which is affected by vehicle holdings (Y_A and Y_B) –

are excluded while forming matched pairs and subsequently during parametric analysis because its inclusion would lead to bad conditioning and induce selection bias (Elwert and Winship 2014).

To create a matched sample following calculation of propensity scores, we use a nearest neighbor matching algorithm which finds 20 or fewer (ratio = 20) control matches for each individual in the treatment group based on propensity score differences. Matching is undertaken with replacement, implying that each control subject may be matched to more than one carshare member. Subjects (both treated and control) without a close enough match in the other group are discarded. The matched subjects are then divided into four subclasses based on their propensity scores such that each subclass has an equal number of members. We chose four subclasses so that we have at least 50 members in each group for statistical reliability. We set a caliper width of 0.2 while undertaking the nearest neighbor matching, i.e. the maximum difference in propensity scores between members of a matched pair is 0.2 standard deviations of the score. This is somewhat more stringent than the 0.25 usually adopted in the literature (Ho et al. 2007). The subsequent balance was deemed adequate and is presented in the next section.

The matched sample, which balances on the confounders listed in Table 1 and thus corrects for bias due to observed variables, is developed to analyze the outcomes listed above.

The matching process was undertaken with replacement of the control units, implying that each control response is matched to a variable number of carsharing member responses. To account for the frequency with which each control unit was matched, weights are estimated by the MatchIt package (Ho et al. 2013). These weights are used in subsequent estimation of summary statistics and parametric modeling.

3.3. Software Packages

Propensity score based matching was undertaken using the MatchIt package in R (Ho et al. 2013). The ordered discrete choice models were estimated using the MASS package in R (Ripley and Venables 2015) and MPLUS (Muthén and Muthén 1998-2014). MPLUS was used for causal mediation analysis (Muthén 2011). Descriptive statistics and graphing were undertaken using the effect package in R (Fox 2003), and MPLUS.

4. Results

4.1. Matched Samples and Covariate Balancing

Table 3 gives the differences in means on key covariates for the original sample ($N_{CS=1}$ = 241, $N_{CS=0}$ = 8,299) and the matched sample ($N_{CS=1}$ = 238, $N_{CS=0}$ = 2298), and the subclasses (subclasses I through IV). The more revealing, but less compact, density plots and 100% stacked column plots for each of the continuous and categorical confounder variables respectively are given in SI-Section 2. In addition to covariate-specific graphs, the SI provides plots giving distributions of the propensity scores of matched and unmatched subjects for all samples.

Table 3. Differences in covariates between carsharing members and non-members (comparison of means)

	Unmatched Sample		**Matched Sample**									
	Entire Sample		**Subclass I**		**Subclass II**		**Subclass III**		**Subclass IV**		**Entire Sample**	
	Members	Control	Members	Control	Members	Control	Members	Control	Members	Control	Members	Control
Sample Size	***241***	***8299***	***60***	***1090***	***60***	***720***	***60***	***342***	***58***	***146***	***238***	***2298***
HH inc. ('000s)												
< $50	8%	14%	13%	12%	5%	8%	10%	8%	4%	5%	8%	9%
$50-100	25%	29%	17%	27%	27%	25%	27%	28%	25%	14%	25%	24%
$100-250	59%	48%	53%	51%	60%	57%	60%	56%	63%	72%	58%	58%
>$250	8%	9%	17%	9%	8%	10%	3%	8%	7%	9%	9%	8%
Residence type												
% single family	53%	83%	93%	90%	72%	67%	35%	39%	10%	15%	54%	54%
Children (#)	0.47	0.56	0.57	0.60	0.55	0.53	0.48	0.44	0.31	0.32	0.47	0.47
HH emp (#)	1.78	1.86	2.0	1.86	1.78	1.80	1.68	1.75	1.76	1.82	1.79	1.79
HH lic (#)	1.92	2.16	2.28	2.17	1.97	2.00	1.77	1.80	1.66	1.80	1.93	1.94
BE (residence)												
Accessibility	74.87	48.59	41.75	44.66	76.82	70.99	84.21	82.61	93.84	92.37	73.82	72.81
Density	1.15	-0.11	-0.36	-0.31	0.92	0.77	1.86	1.79	2.17	2.23	1.14	1.10
Commute dist (miles)	11.8	14.9	15.90	15.43	14.82	14.05	10.98	12.56	6.01	7.03	11.9	14.1
BE (work place)												
Accessibility	76.85	56.75	57.75	54.65	74.45	75.12	84.15	82.68	91.36	91.24	76.79	75.81
Density	1.49	0.57	0.60	0.52	1.33	1.35	1.86	1.78	1.99	2.27	1.49	1.46
Age (years)	45.72	48.54	48.30	48.91	46.80	46.29	46.65	44.89	41.49	44.63	45.83	45.99
Male	61%	51%	63%	49%	53%	60%	57%	65%	66%	71%	61%	59%
Undergrad	41%	31%	32%	33%	38%	41%	48%	41%	44%	43%	40%	39%
Graduate	46%	32%	45%	33%	44%	45%	44%	53%	54%	57%	47%	45%

The four subclasses in the matched sample represent distinct groups with different propensities for joining carsharing. On one end, subclass I consists of respondents (both members and non-members) who live and work in suburban neighborhoods, reside in single family homes, have adequate access to vehicles (at least one car per licensed driver in HH), have long commuting distances, and largely drive or carpool to work. These subjects have a limited propensity to join carsharing (1-2% probability). On the other end, subclass IV comprises respondents who live and work in city centers, have limited access to personal vehicles, largely dwell in multi-unit houses, do not drive to work, have a reasonably short commute distance, and almost certainly have a college degree. These subjects are somewhat younger compared to the members in Subclass-I (44 years versus 48), and have a relatively high propensity to join carsharing (~40%).

Subclasses II and III have characteristics that fall in between these two extremes. For example, as we move from I to IV, a smaller proportion of respondents lives in single family homes. Similarly, the share of car as a primary commute mode as well as the average commute distance decreases. There is also a decrease in the average number of children in the household and an increase in the proportion of respondents with at least an undergraduate degree. However, the built environments of both residential and workplace neighborhoods of II and III are close to IV in terms of density and accessibility and hence may be considered urban. In other words, three quarters of the carsharing respondents in the matched sample live as well as work in urban neighborhoods.

Comparison of the average covariate values (Table 3 and SI Section S2), as well as plots showing the distributions of the estimated propensity score between treated and untreated subjects in the matched samples, indicates that, while the covariates are well-balanced between the two groups for the matched sample as a whole, and while most of the covariates are well-balanced for the subsamples as well, some differences persist at the subsample level. For example, 10% of carshare members in subclass IV live in single-family homes, compared to 15% of non-members. Tightening the caliper below 0.15 would have led to a further reduction in matched subjects in subclass IV. To control for these within-class differences, we conducted parametric analysis (Section 4.3) after inspection of simple (weighted) differences in means and medians of the outcome variables (Section 4.2).

4.2. Summary Statistics of Outcomes

SI Section 4 gives the proportion of respondents in various levels of vehicle holdings (both total household vehicles and vehicles per driver). For the entire matched sample, one-third of members are from households without any vehicles, and two-thirds are from households with less than one car per driver (including zero). This may be compared to around 62% members from households without vehicles found by Martin, Shaheen, Lidicker (2010) for all of North America. The corresponding proportions for non-members are 11% and 43% respectively. We should point out that these differences may not be entirely be attributed to differences in carsharing membership status because, as a result of selecting "nearest" rather than "exactly" matched control units, we have not completely controlled for the differences in confounders.

SI Section 4 also disaggregates the responses by subclasses, which reflect the propensity to join carsharing. In subclass IV, which reflects the highest propensity to enroll in carsharing, two-thirds of members are carless and nearly all members are from households with more drivers than cars. These ratios decrease monotonically as we move down from subclass IV to I. For weekly transit and non-motorized trips, members are more likely to be in higher usage classes than non-members within any subclass (Section 4.3).

Considering the original "continuous" variables, we find that household vehicle holdings per driver decrease monotonically from subclass I to IV for both members and non-members.

Correspondingly, there is a monotonic increase in weekly one-way transit and non-motorized trips (except for a small dip in transit trips from III to IV). On average, respondents in subclass I have 50% more vehicles per license holder (Δ = 0.3), and take half as many transit (Δ = 2.5) and non-motorized trips (Δ = 5) as those in IV. The difference in vehicle holdings between members and non-members increases as we move from subclass I to IV. In subclass I, both members and non-members have access to at least one vehicle, while in subclass IV the non-members have three times greater access than members (Δ = 0.4). Across most subclasses, carshare members make more transit, bike, and walk trips than non-members.

4.3. Differences in vehicle holdings

Table 4 gives the summary of the two ordered probit models estimated for household vehicles and household vehicles per driver. The first two rows of Table 4A give the CS coefficients of unstandardized and Y-standardized models respectively where the latter is obtained by dividing the estimated parameters from the unstandardized model by the standard deviation of the dependent variable U* (Long 1997 p. 169). The coefficients are negative in sign across all subclasses, indicating that carsharing members in all groups have lower vehicle holdings - both on a total basis and on a per-driver basis. The Y-standardized coefficients may be interpreted as the difference in latent propensities to own more vehicles (U*) between carsharing members and non-members in standard deviation (SD) units. For urban members (subclasses II to IV), propensity to own more vehicles is lower by 0.3 to 1.3 SD units; and utility of holdings per licensed driver is lower by 0.5 to 1.8 SD units. These differences are both substantial and statistically significant. On the other hand, for suburban dwellers, the differences are significant and moderately substantial (0.3 SD units) for total household vehicles, but small and insignificant for vehicles per driver.

Table 4A. Models for vehicle holdings – coefficients of carsharing membership indicator

		Subclass I	Subclass II	Subclass III	Subclass IV
U*: Utility from Total Household Vehicle Holdings					
Coefficients of CS	Unstandardized	**-0.45**	**-0.49**	**-1.03**	**-2.03**
(Member=1)	Y - Standardized	**-0.30**	**-0.33**	**-0.69**	**-1.32**

Marginal Effects of CS membership	(0) Prob ($Y_1 = 0$)	0.02	**0.02**	**0.07**	**0.29**
	(1) Prob ($Y_1 = 1$)	0.17	**0.18**	**0.31**	**0.33**
	(2) Prob ($Y_1 = 2$)	-0.15	**-0.16**	**-0.33**	**-0.55**
	(3) Prob ($Y_1 \geq 3$)	-0.04	**-0.05**	**-0.06**	**-0.07**
U*: Utility from Vehicle Holdings per Licensed Driver					
Coefficients of CS (1 = Member)	Unstandardized	-0.17	**-0.55**	**-0.94**	**-2.35**
	Y - Standardized	-0.15	**-0.49**	**-0.84**	**-1.79**
Marginal Effects of CS membership	(0) Prob ($Y_{1A} < 1$)	0.08	**0.22**	**0.36**	**0.61**
	(1) Prob ($Y_{1A} = 1$)	-0.06	**-0.17**	**-0.31**	**-0.55**
	(2) Prob ($Y_{1A} > 1$)	-0.02	**-0.04**	**-0.05**	**-0.06**

Table 4B. Models for vehicle holdings – coefficients of covariates

	# of HH Vehicles		**Vehicles per Driver**	
	Un-Standardized	**Y - Standardized**	**Un-Standardized**	**Y - Standardized**
Household Variables				
Density - Residence	**-0.12**	**-0.08**	**-0.16**	**-0.12**
Density - Workplace	**-0.07**	**-0.05**	**-0.06**	-0.05
Annual HH Income (Ref: > $250K)				
< $100K	**-0.50**	**-0.34**	**-0.46**	**-0.36**
$100 - $150K	**-0.34**	**-0.23**	**-0.31**	**-0.23**
$150 - $250K	0.15	0.10	-0.20	-0.15
# of children in HH	0.07	0.05	0.04	0.02
# of drivers in HH	**1.23**	**0.84**	**-0.49**	**-0.37**
Individual Variables				
Education (Ref: Graduate degree)				
Some college or less	**0.32**	**0.22**	**0.44**	**0.33**
Undergraduate degree	**0.19**	**0.13**	**0.26**	0.19
Age	0.01	0.00	0.00	0.00
Commute Distance (log)	0.18	0.12	-0.20	0.15
Gender (Ref. Male)	-0.01	-0.01	-0.05	-0.04
Thresholds				
(0) / (1)	-0.35	-0.23	**-1.93**	**-1.46**
(1) / (2)	**1.64**	**1.08**	0.40	0.30
(2) / (3)	**3.75**	**2.45**		
R-squared	0.58		0.22	

Notes: (1) Figures in bold are statistically significant at α = 0.05 or lower.

The marginal effects of carsharing indicate that within any subclass, members tend to have a lower utility than non-members for owning some number of vehicles, after conditioning for various confounding variables (Table 4B). We should reiterate that the above results may be considered to hold in the sub-population whose observed characteristics match those of carsharing members;

these results may not be generalized to the larger population of employed residents of the San Francisco Bay Area. In other words, the differences may be interpreted as the average treatment effect on the treated, or ATT.

4.4. Effect on frequency of transit and walk and bike trips

Table 5 summarizes the differences between carsharing members and non-members in their standardized frequency propensity (variance = 1) for usage of transit and non-motorized modes. The differences are disaggregated into (a) *direct* differences after conditioning on all confounders including vehicle holdings, (d) *mediated* differences resulting from differences in vehicle holdings (HHVeh), and (c) *total* differences which sums the above two and may be interpreted as differences after conditioning on all confounders *except* vehicle holdings.

The positive sign of the coefficients indicates that the propensity to make more non-auto trips is higher for carsharing members than for non-members within any subclass. Between the two groups, the propensity to use transit more often is higher by 0.13 to 0.24 standard deviations. The propensity to walk and bike more often is higher by 0.17 to 0.53 standard deviations.

Ex-ante, we expected the direct pathways to be negative, assuming that carsharing provides additional opportunity for driving to a destination and thus potentially substitutes for transit and non-motorized modes. However, nearly all the coefficients are positive, albeit statistically not significant implying that we cannot reject the hypothesis that there are no differences between members and non-members on these propensities, after controlling for confounding variables.

Table 5: Models for frequency of use of transit and non-motorized modes

		Subclass I	**Subclass II**	**Subclass III**	**Subclass IV**
Latent indicator of transit frequency (V_C^* & W_C^*)					
Coefficients of CS	- Direct pathway	0.24	0.10	0.05	0.11
(Member=1)	- Mediated pathway(s)	0.00	**0.04**	**0.11**	0.07
	- Total	0.24	0.13	0.17	0.19
Latent indicator of frequency to bike and walk (V_D^* & W_D^*)					
Coefficients of CS	- Direct pathway	**0.36**	0.10	**0.37**	**0.41**
(Member=1)	- Mediated pathway(s)	0.04	**0.07**	**0.16**	-0.01
	- Total	**0.41**	0.17	**0.53**	**0.40**

Notes: (1) All coefficients reported in this table are Y-Standardized. (2) Figures in bold are statistically significant at α = 0.05 or lower.

5. Discussion and conclusions

We find that vehicle holdings of carsharing members are substantially and significantly lower than those of non-members with similar characteristics in terms of individual and household demographics as well as built environment features of both residential and job location. The differences increase with increases in propensity to enroll for carsharing. The latent propensity to own more vehicles is lower for members by 0.3 to 1.3 standard deviations relative to non-members. When measured on a per household driver basis, the propensity is lower by 0.2 to 1.8 standard deviations. Members are also likely to walk and bike, and use transit more frequently than non-members. However, these differences are relatively minor and often not statistically significant.

Our examination of the causal impacts of carsharing on travel behavior utilizing the California Household Travel Survey (CHTS) has several limitations pertaining to the cross-sectional nature of the data and the methodology adopted that does not control for all biases in the data.

(a) Limitations of cross-sectional data:

Causal relationships imply a temporal sequence, with the treatment (carsharing enrollment) preceding the outcome (change in travel behavior). For mediated analysis, a drop in vehicle holdings precedes changes in other aspects of travel behavior. A limitation of cross-sectional datasets such as the CHTS is that they capture the state of affairs at a snapshot of time and hence are not as definitive as a longitudinal study in establishing causal relationships (Kline 2011 p. 98). Further, there is an implicit assumption of equilibrium, meaning that any causal impacts (travel behavioral changes) have manifested and the system is in a steady state. In other words, the causal process has "…dampened out…" (Kline 2011 pp 121). From the available dataset, it is not feasible to assess the stability of the system. Further, it is likely that the membership duration varies across respondents – in which case causal effects are at different stages of manifestation. Our analysis disregards these dynamic processes.

(b) Not all sources of selection biases have been addressed:

Matching does not correct for selection bias arising due to unobservables – variables considered by the decision maker but not observable by analysts (Heckman et al. 1996; Caliendo and Kopeinig 2008). In this analysis the key unobservables are attitudes of the decision makers – which have been shown to have a large effect on travel and residential location decisions (Handy, Cao and Mokhtarian 2005). Unfortunately, the CHTS dataset does not enable us to directly gauge the attitudes of individuals. How large is the bias resulting from these unobserved variables? To the extent that unobservables are highly correlated with observed covariates that we balance for during matching, we can assume that this bias is small. For example, individuals with pro-bike and pro-walk attitudes are likely to live in urban neighborhoods, while individuals with car dependent attitudes are likely to drive for their daily commute. By balancing on built environment characteristics of residential neighborhoods as well as primary commute mode, we can argue that some of the selection bias that arises due to these unobserved attitudes has been balanced.

A number of other methods in the statistical and econometric literature also seek to correct for selection bias in observational and cross-sectional data: instrumental variable(s) models, selection models, joint discrete choice models, and structural equation models (Mokhtarian and Cao 2008). Various selection models like the Heckman Two-Step (or Heckit) estimator are designed to explicitly account for unobserved variables driving decision-making, but require various exclusion restrictions and are sensitive to model specification (Briggs 2004; Bushway, Johnson and Slocum 2007).

(c) Reverse causality needs to be addressed:

Perhaps, the biggest shortcoming in our model is that we do not account for the simultaneity or reverse causality bias – specifically the fact that vehicle ownership levels are not only influenced by carsharing membership but also influence the decision to join carsharing. Largely because of this, we cannot give a causal inference to the observed differences even though we may have successfully controlled for selection bias. In their survey of carsharing members, Martin, Shaheen, Lidicker (2010) asked the respondents about the circumstances under which they joined carsharing. Around 30% indicated that they joined carsharing to avoid purchasing an additional vehicle or to dispose of their vehicle, highlighting the impact of carsharing on vehicle ownership. Further, around 50% indicated that they did not own a vehicle and joined carsharing to get access to vehicles, highlighting the strength of the relation in the reverse direction.

In the context of a parametric model (such as regression) to explain the influence of carsharing on travel behavior, both simultaneity bias and selection bias may be framed as endogeneity problems. Endogeneity arises when one or more of the explanatory variables are correlated with error term ε leading to biased (incorrect) estimates of the regression coefficients for treatment indicator CS and X, and their standard errors. Such correlation arises when important variables like attitudes towards driving and car ownership, influencing both propensity to join carsharing (CS) and travel behavior outcomes (Y), are omitted (selection bias due to unobserved differences in characteristics). It also arises when propensity to join carsharing not only affects Y, but is also affected by Y. In a subsequent paper, we will explore alternative statistical methods to jointly control for both selection and simultaneity biases.

The CHTS survey asked respondents about their membership status but collected no information about the duration of membership, frequency of carsharing usage, and conditions under which a respondent joined carsharing. Sioui, Morency, Trépanier (2013) divided their sample into four groups based on frequency of carsharing usage, and found that the modal share for daily trips varied substantially between the groups after controlling for household size and level of car ownership. Given the growing prevalence of shared-use mobility and its potential to reduce car dependency in urban regions, we recommend that future large-scale regional travel surveys such as the CHTS and National Household Travel Survey (NHTS) collect these additional details.

This analysis of the impacts of carsharing on travel behavior suggests that after accounting for the complex confounding factors that also influence travel behavior, some effects remain. Urban carshare members still tend to engage in more sustainable travel patterns (lower vehicle holdings and increased transit, bike, and walk trips) than non-carshare members, after partially accounting for self-selection bias. Although significant research remains to be done, this study provides additional evidence that the availability of carsharing services may enable more sustainable travel behavior.

6. Acknowledgements

The study benefited from input and critical comments from various researchers including Kenneth Laberteaux from Toyota Research Institute of North America; and Susan Handy, Gil Tal, Sonia Yeh, Lewis Fulton, David Bunch and Joan Ogden, all from Institute of Transportation Studies at UC Davis. We are also grateful to Karim Hamza for help with obtaining accessibility scores. The views and opinions expressed in this paper are those of the authors alone.

7. References

Angrist, J., Pischke, J.-S.: Mostly harmless econometrics: An empiricist's companion. Princeton University Press, (2008)

Austin, P.C.: An introduction to propensity score methods for reducing the effects of confounding in observational studies. Multivariate Behavioral Research **46**(3), 399-424 (2011). doi:DOI:10.1080/00273171.2011.568786

Bhat, C.R., Eluru, N.: A Copula-based approach to accommodate residential self-selection effects in travel behavior modeling. Transportation Research Part B: Methodological **43**(7), 749-765 (2009)

Bhat, C.R., Pulugurta, V.: A Comparison of Two Alternative Behavioral Choice Mechanisms For Household Auto Ownership Decisions. Transportation Research Part B: Methodological **32**(1), 61-75 (1998)

Briggs, D.C.: Causal inference and the Heckman model. Journal of Educational and Behavioral Statistics **29**(4), 397-420 (2004)

Brookhart, M.A., Schneeweiss, S., Rothman, K.J., Glynn, R.J., Avorn, J., Stürmer, T.: Variable Selection for Propensity Score Models. American Journal of Epidemiology **163**(12), 1149-1156 (2006)

Bushway, S., Johnson, B.D., Slocum, L.A.: Is the Magic Still There? The Use of the Heckman Two-Step Correction for Selection Bias in Criminology. Journal of Quantitative Criminology **23**(2), 151-178 (2007)

Caliendo, M., Kopeinig, S.: Some Practical Guidance for the Implementation of Propensity Score Matching. Journal of Economic Surveys **22**(1), 31-72 (2008)

Caltrans: 2010-2012 California Household Travel Survey Final Report. In. California Department of Transportation, (2013)

Cervero, R., Golub, A., Nee, B.: City CarShare: Longer-Term Travel Demand and Car Ownership Impacts. Transportation Research Record: Journal of the Transportation Research Board **1992**(1), 70-80 (2007)

Cervero, R., Tsai, Y.: City CarShare in San Francisco, California: Second-year Travel Demand and Car Ownership Impacts. Transportation Research Record: Journal of the Transportation Research Board **1887**(1), 117-127 (2004)

Elwert, F., Winship, C.: Endogenous Selection Bias: The problem of conditioning on a collider variable. Annual Review of Sociology **40**, 31-53 (2014). doi:10.1146/annurev-soc-071913-043455

Firnkorn, J., Müller, M.: What will be the Environmental Effects of New Free-Floating Car-Sharing Systems? The Case of Car2go in Ulm. Ecological Economics **70**(8), 1519-1528 (2011)

Fox, J.: Effect Displays in R for Generalised Linear Models. Journal of Statistical Software **8**(15), 1-27 (2003)

Greene, W.H., Hensher, D.A.: Modeling Ordered Choices - A Primer. Cambridge University Press, Cambridge, UK (2010)

Handy, S., Cao, X., Mokhtarian, P.: Correlation or causality between the built environment and travel behavior? Evidence from Northern California. Transportation Research Part D: Transport and Environment **10**(6), 427-444 (2005). doi:10.1016/j.trd.2005.05.002

Heckman, J.J., Ichimura, H., Smith, J., Todd, P.: Sources of Selection Bias in evaluating social programs: An interpretation of conventional measures and evidence on the effectiveness of Matching as a program evaluation metho. Proceedings of the National Academy of Sciences **93**(23), 13416-13420 (1996)

Ho, D.E., Imai, K., King, G., Stuart, E.A.: Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. Political Analysis **15**(3), 199-236 (2007)

Ho, D.E., Imai, K., King, G., Stuart, E.A.: MatchIt: Nonparametric Preprocessing for Parametric Causal Inference - R Package Version 2.4-21. In. (2013)

Ho, D.E., Rubin, D.B.: Credible Causal Inference for Empirical Legal Studies. Annual Review of Law and Social Science **7**, 17-40 (2011)

Imai, K., Keele, L., Tingley, D.: A General Approach to Causal Mediation Analysis. Psychological Methods **15**(4), 309 (2010)

Imbens, G.W., Rubin, D.B.: Causal Inference in Statistics, Social, and Biomedical Sciences. Cambridge University Press, (2015) Klincevicius, M.G., Morency, C., Trépanier, M.: Assessing Impact of Carsharing on Household Car Ownership in Montreal, Quebec, Canada. Transportation Research Record: Journal of the Transportation Research Board **2416**(1), 48-55 (2014)

Kline, R.B.: Principles and Practice of Structural Equation Modeling, Third ed. Guilford press, New York, NY (2011)

Long, J.S.: Regression Models for Categorical and Limited Dependent Variables. SAGE, (1997)

Loose, W.: The State of European Car-Sharing. In., p. 129. Bundesverband CarSharing e. V., (2010)

MacKinnon, D.P., Fairchild, A.J., Fritz, M.S.: Mediation analysis. Annual review of psychology **58**, 593 (2007)

Martin, E., Shaheen, S.: The impact of carsharing on public transit and non-motorized travel: An exploration of North American carsharing survey data. Energies **4**(11), 2094-2114 (2011a). doi:10.3390/en4112094

Martin, E., Shaheen, S.A., Lidicker, J.: Impact of carsharing on household vehicle holdings. Transportation Research Record: Journal of the Transportation Research Board **2143**(1), 150-158 (2010). doi:10.3141/2143-19

Martin, E.W., Shaheen, S.A.: Greenhouse Gas Emission Impacts of Carsharing in North America. IEEE Transactions on Intelligent Transportation Systems **12**(4), 1074-1086 (2011b)

Mokhtarian, P.L., Cao, X.: Examining the Impacts of Residential Self-selection on Travel Behavior: A Focus on Methodologies. Transportation Research Part B: Methodological **42**(3), 204-228 (2008). doi:DOI: 10.1016/j.trb.2007.07.006

Moulton, B.R.: An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units. The Review of Economics and Statistics, 334-338 (1990)

Muthén, B.: Applications of Causally Defined Direct and Indirect Effects in Mediation Analysis using SEM in Mplus. (2011)

Muthén, L.K., Muthén, B.O.: Mplus User's Guide Seventh Edition. In. Muthén & Muthén, Los Angeles, CA, (1998-2014)

Ripley, B., Venables, W.: MASS (Modern Applied Statistics with S) - R Package Version 7.3-42. In. (2015)

Rosenbaum, P.R., Rubin, D.B.: The Central Role of the Propensity Score in Observational Studies for Causal Effects. Biometrika **70**(1), 41-55 (1983)

Shaheen, S.A., Cohen, A.P., Chung, M.S.: North American Carsharing - 10-Year Retrospective. Transportation Research Record: Journal of the Transportation Research Board **2110**(1), 35-44 (2009)

Sioui, L., Morency, C., Trépanier, M.: How Carsharing Affects the Travel Behavior of Households: A Case Study of Montréal, Canada. International Journal of Sustainable Transportation **7**(1), 52-69 (2013). doi:DOI:10.1080/15568318.2012.660109

Stillwater, T., Mokhtarian, P.L., Shaheen, S.A.: Carsharing and the built environment. Transportation Research Record: Journal of the Transportation Research Board **2110**(1), 27-34 (2009)

Tal, G.: Evaluating the Effect of Car-Sharing: Exploring the Gap Between What We Know vs. What We Need to know and Its Effect on Optimism Bias. In., p. 32. Institute of Transportation Studies, University of California, Davis, (2009)

U.S. Census Bureau: American FactFinder - 2013 ACS 3-year estimates. U.S. Department of Commerce. http://factfinder.census.gov/ (2014).

Weitzen, S., Lapane, K.L., Toledano, A.Y., Hume, A.L., Mor, V.: Principles for Modeling Propensity Scores in Medical Research: A Systematic Literature Review. Pharmacoepidemiology and Drug Safety **13**(12), 841-853 (2004)

Zhou, B., Kockelman, K.M.: Opportunities for and Impacts of Carsharing: A Survey of the Austin, Texas Market. International Journal of Sustainable Transportation **5**(3), 135-152 (2011)

PART-III

Paper 2: Addressing the joint occurrence of self-selection and simultaneity biases in the estimation of program effects based on cross-sectional observational surveys - Case study of travel behavior effects in carsharing

Resubmitted for review to Transportation (Springer)

ADDRESSING THE JOINT OCCURRENCE OF SELF-SELECTION AND SIMULTANEITY BIASES IN THE ESTIMATION OF PROGRAM EFFECTS BASED ON CROSS-SECTIONAL OBSERVATIONAL SURVEYS - CASE STUDY OF TRAVEL BEHAVIOR EFFECTS IN CARSHARING

Gouri Shankar Mishra (Corresponding author)
Institute of Transportation Studies, University of California, Davis
gsmishra@ucdavis.edu

Patricia L. Mokhtarian
School of Civil and Environmental Engineering, Georgia Institute of Technology
patmokh@gatech.edu

Regina R. Clewlow
Institute of Transportation Studies, University of California, Davis
rclewlow@ucdavis.edu

Keith F. Widaman
Graduate School of Education, University of California, Riverside
keith.widaman@ucr.edu

Abstract

We estimate the effect of carsharing on travel behavior (specifically, household vehicle holdings, frequency of transit usage, and frequency of biking and walking) using data from the 2011-12 California Household Travel Survey (CHTS). The effect of carsharing on vehicle ownership is a dynamic process that plays out over a period of time – past ownership influences enrollment decisions, which in turn influence holdings in a later period. Representing this process using cross-sectional data conflates causal effects with simultaneity bias. Further, members and non-members differ in various observed and unobserved ways - demographics, built environment of residential and workplace location, and attitudes - raising the potential for self-selection bias in comparing travel behavior between the two groups given the observational nature of the data.

Drawing on established methods for dealing with each bias individually, we develop a method to help control for this joint occurrence of self-selection and simultaneity biases. Restricting the analysis to employed respondents residing in the San Francisco Bay Area, we find that 80% of the observed difference of 0.9 units in average vehicle holdings between carsharing non-members and members may be explained by the biases listed above. The remaining difference of 0.17 units reflects the estimated effect of carsharing, which is the equivalent of shedding one vehicle by about one out of every six households whose member(s) are enrolled in carsharing. The effect on transit usage and walking and biking frequency is positive, albeit small and statistically non-significant. We identify factors that may affect the internal and external validity of our results. Our methods cannot completely replace randomized experiments or panel data. However, the methods used here provide a way to help control for the joint occurrence of self-selection and simultaneity biases, and provide a ballpark estimate of causal effects, for large-scale, general-purpose, cross-sectional datasets such as the CHTS.

Keywords: *Future mobility, program evaluation, omitted variable bias, treatment effects, vehicle ownership, selection on unobservables*

1. Introduction

It is often stated that the gold standard for estimating the effects of a treatment or program is a longitudinal study with randomized assignment of respondents to the program or no-program (control) groups (Rubin 2008). Random assignment reduces the chance of systematic differences between the treatment and control groups before program assignment and thus (to the extent that it does result in statistically equivalent groups) controls for selection bias. Longitudinal data confirm the temporal sequence of events – outcome (usually) follows the treatment, allowing analysts to distinguish the effect of program membership on outcome from the effect of outcome status on propensity to join (or remain in) the program. Unfortunately, many evaluations are based on cross-sectional observational studies, and therefore possess the potential for both selection bias and simultaneity bias.

Observational study participants may have selected themselves into the program under consideration – hence naïve comparison of outcomes between treated and control members reflects not only the effect of the program under study, but also selection bias. Selection bias may arise from differences in characteristics between treatment and control members. The characteristics of interest are those that influence the decision to join the program, as well as the behavioral response to the program. These characteristics could include individual and household demographics, attitudes towards issues (such as the environment or adoption of new technology), employer policies and social norms. When these characteristics are observed, selection bias may arise in two possible ways (Heckman et al. 1996). Selection bias may arise due to non-overlapping observed characteristics (B1; see Table 1), i.e. observations outside the common support, in the parlance of the treatment effects literature. In other words, this bias occurs when for some control members there are no comparable program participants with similar characteristics, and/or vice versa. Additionally, selection bias results from differences in distribution of observed characteristics within the common support (B2). Differences in distribution refer to differences in expectation (average value), in dispersion (usually referring to variance or its square root, standard deviation), and in shape (usually summarized by third and fourth moments, i.e. skewness and kurtosis).

If the observed differences are assumed to entirely account for the selection bias, then the situation is variously referred to in the literature as unconfoundedness, selection on observables, exogeneity, and conditional independence (Elwert and Winship 2014). A more reasonable assumption is that not all variables affecting program participation or selection are observed by analysts or researchers, yielding selection on unobservables. Such a situation leads to an omitted variable bias (B3) (Angrist and Pischke 2008 pp. 44-47; Elwert and Winship 2014; Heckman et al. 1996).

In addition to self-selection bias, studies estimating causal effects of programs from cross-sectional data may have to confront the potential for reverse causality bias or simultaneity bias (B4). Simultaneity bias is produced when an "explanatory" variable is simultaneously a function of the "dependent" variable it is supposed to explain – that is, when one variable is both a cause and an effect of another (Mokhtarian and Cao 2008).

In addition to the biases B1 to B4 listed above, we researchers may induce bias by bad conditioning (B5), i.e. by controlling for a variable that represents an outcome of both the program participation and outcome variables being studied (Elwert and Winship 2014; Angrist and Pischke 2008 pp. 47-49). By "controlling", we mean for example incorrectly including the variable as an explanatory variable in a parametric (example: regression) analysis. Or we may incorrectly restrict our analysis to a subsample formed on the basis of the variable (for example, we may only consider a sample of urban dwelling respondents).

In this study, we develop a framework to make causal inferences from cross-sectional data when there is potential for both self-selection and simultaneity biases. The framework uses a parametric analysis where propensity to participate in the program and its effect on outcome are simultaneously estimated in a structural equation model (SEM). The parametric model allows for feedback loops between participation indicator and outcome (referred to as a non-recursive model), to estimate the influences in both directions and control for simultaneity bias (B4). Further, the residual errors or disturbance terms of the endogenous variables are allowed to co-vary, capturing the aggregated confounding influence of all omitted variables and thus controlling for omitted variable bias (B3) (Kline 2012 p. 107). The SEM is estimated on a matched sample where one or more control units are matched to each treatment unit such that the units have the same or similar values for various observed covariates, thereby minimizing any selection bias arising because of differences in observed characteristics (B1 and B2). We judiciously exclude variables influenced by the outcomes being studied from both the matching process and subsequent parametric analysis to avoid bad conditioning (B5).

We apply the approach in a travel behavior context, in which we estimate the effect of carsharing membership on several measures of travel behavior, namely vehicle holdings and frequency of transit usage and biking and walking, using data from the 2010-2012 California Household Travel Survey or CHTS (Caltrans 2013). Exploiting the same dataset, Mishra et al. (2015) controlled for self-selection bias arising due to *observed* differences in characteristics between carsharing members and non-members, using propensity score matching. Each carsharing member was matched to non-members who have similar individual and household demographics and reside in neighborhoods with similar built environments. However, the previous study cannot claim that carsharing caused the observed differences in travel behavior between the matched pairs, due to the presence of simultaneity bias as well as likely self-selection bias from differences in *unobserved* characteristics. In this paper, we implement an approach to help control for these biases, to improve our ability to estimate the effects of carsharing on travel behavior.

2. Carsharing and travel behavior

Since the arrival of Zipcar™ in 2005, carsharing has grown in popularity in major metropolitan areas around the world, as well as evolved into new, more convenient models, representing a potential shift from vehicles as an owned asset to vehicles as a mobility service. This paper focuses on *traditional*, or station-based carsharing services, the model which served the majority of users when the California Household Travel Survey was deployed. In station-based carsharing services, users often join through a membership model, providing them with access to borrow vehicles from designated locations on an hourly basis. Through this model, the vehicles must be returned to the same location, unlike newer forms of *free-floating* carsharing systems that allow the pick-up and return of vehicles in different locations and often charge on a per-minute basis (e.g., car2go®, DriveNow®, and GoDrive®). Another related shared mobility service model is based on ride-hailing services, through which users can pay for rides in private vehicles (e.g. Uber®, Didi Kuaidi®). Although there is an emerging body of literature on the potential impacts of newer shared mobility services, much of what is understood about the impacts of carsharing on travel behavior is based on studies of traditional, station-based carsharing.

Self-selection bias in the context of carsharing may be framed in the following way: people who self-select to become members may differ from non-members in terms of socio-economic and demographic characteristics, residential location choices, and attitudes. Our analysis of the CHTS highlighted large differences in observed characteristics between members and non-members, as summarized in Table 3. Members are more likely to reside in dense urban neighborhoods where carsharing is available, unlike non-members, who predominantly reside in suburban regions. Focusing only on employed respondents in the Bay Area, we find that members are more likely to be male (61% versus 51%), highly educated (87% with a college degree or higher versus 64%), and reside in multi-unit housing complexes rather than single family homes (47% versus 17).These observations are consistent with those of other researchers using different datasets (Martin et al. 2010; Sioui et al. 2013; Loose 2010).

The potential for omitted variable bias arises because information on important confounding variables may not be collected as part of a survey. For example, large regional travel surveys like the CHTS and National Household Travel Survey (NHTS) measure a wide range of demographic characteristics but collect little information on attitudes towards such issues as the environment, driving, transit use, and urban living. However, such attitudes have been shown to considerably improve the ability to explain people's travel and residential location decisions (Handy et al. 2005).

We should note that the omitted variable bias of main concern to the present study arises when the unobserved variable is correlated with both endogenous variables of interest: the outcome variable and the variable indicating program participation. For example, consider an employer-sponsored travel demand management (TDM) scheme whereby members receive subsidized carsharing membership if they commute to work by transit or bike or walk. Here the subsidy influences not just program participation but also other travel behavior outcomes. In other words, in the absence of information about the TDM scheme, the observed correlation may be incorrectly attributed to the effect of carsharing on transit usage, leading to omitted variable bias. Another example of an omitted variable would be a positive attitude towards driving that may positively affect both vehicle ownership and carsharing enrollment. Although less a focus of attention in the present study, omitted variable bias can also arise from correlation of an unobserved variable with only one of the endogenous variables, and with one or more exogenous variables. For example, an unobserved status-seeking attitude may influence both residential location type (exogenous) and vehicle ownership (endogenous). Not controlling for such omitted variables will influence the estimated effect of the residential location type on vehicle ownership; however, identifying such an effect is not the primary objective of the current study.

Table 1. Potential biases while drawing causal inferences from cross-sectional observed data, using examples from carsharing based on the 2010-12 CHTS dataset.

Bias	Source	Carsharing example	Method adopted for control of bias
B1	**Non-overlapping support:** Non-overlapping observed characteristics of program participants and non-participants.+	Around 2% of the sample of employed people in the San Francisco (SF) Bay Area do not have a driver's license. These non-members do not have any equivalent members.	Drop responses with propensity scores outside the range of common support.
B2	**Differences within support**: Differences in distribution of observed characteristics within the common support.+	On average, members are more educated: 87% have an undergraduate degree or higher compared to 63% of non-members. Similarly, members live predominantly in urban areas, while non-members are predominantly suburban dwellers.	Matched sampling ensures that probability distribution of X among controls in matched sample is close to that of members. Parametric analysis on matched sample controls for residual differences.
B3	Selection on unobservables or **omitted variable bias**: Unobserved variables influencing both the decision to sign up for treatment and the behavioral outcome.	People with a general dislike for driving are likely to own fewer vehicles in the household and unlikely to choose carsharing over transit for their mobility requirements. However, there is no information in the CHTS to gauge attitudes towards driving	Path model assumes a non-zero correlation between the disturbances (errors) of endogenous variables, capturing association due to common unobserved / omitted influences on those variables.
B4	**Reverse causality bias or simultaneity bias**: Program participation may be caused by the variable *labeled as* the outcome, suggesting a y➔d causal pathway.	Vehicle holdings (*y*) influence the propensity to enroll in carsharing (*d*). Carsharing (*d*), on the other hand, incentivizes reduction in vehicle holdings (*y*) (our initial hypothesis).	Non-recursive path model with unidirectional arcs between vehicle holdings and the carsharing membership indicator (feedback loops).
B5	**Bad conditioning**: Conditioning on a variable that represents an outcome of both the program participation and outcome variable(s) being studied (++)	Commute mode choice is known to depend upon vehicle holdings, and is likely to depend on carsharing participation. Hence, a regression of vehicle holdings on carsharing and various socio-demographic variables with commute mode choice as a covariate will bias the coefficient estimates.	Judicious exclusion of key variables from propensity score model adopted to identify a matched sample, as well as from the subsequent parametric model.

Notes: (+) The examples in this table are based on differences between members and non-members on a single covariate. Similar logic applies in the k-dimensional covariate space where k represents the number of covariates, their interactions, and second-order terms as well as other transformations such as log and exponential. (++) Elwert, Winship (2014) refer to this bias as an endogenous selection bias, while referring to B1-B3 as examples of "confounding variable bias".

The potential for simultaneity bias in the context of the effect of carsharing on vehicle ownership based on cross-sectional data may be framed in the following way: carsharing members indeed have the incentive to reduce their vehicle holdings given the large costs of ownership in urban neighborhoods (e.g., parking), but simultaneously, people without adequate vehicle access join carsharing to enhance their mobility options (vehicle ownership influences carsharing adoption). Using a carefully designed survey administered to 9000+ members across the U.S. and Canada in 2008, Martin et al. (2010) highlighted the complex two-way relationship between carsharing and vehicle holdings. The study asked the respondents about the circumstances under which they joined carsharing. Around 30% indicated that they joined carsharing to avoid purchasing an additional vehicle or to dispose of their vehicle, highlighting the impact of carsharing on vehicle ownership. Even more, however (around 50%) indicated that they did not own a vehicle and joined carsharing to get access to vehicles - highlighting the strength of the relation in the reverse direction.

Much of the research on the causal effects of carsharing is based on surveys of carsharing members using purpose-built survey instruments. Some of these studies address the dynamic two-way relationship between carsharing and vehicle holdings, and thus control for reverse causality bias by asking respondents about the number of vehicles held before and after enrolling in carsharing; e.g. Firnkorn, Müller (2012) and Martin et al. (2010). Further, the issue of self-selection bias takes a different form than that addressed in this paper, given that these member-only studies assess effects not by comparing changes to a control group, but by letting respondents characterize their decision behind vehicle disposal and carsharing enrollment. For example, in the survey by Firnkorn, Müller (2012), members of the car2go® service in Ulm, Germany attribute their decisions to dispose or forgo future purchases of a vehicle either to the car2go® service or to reasons unrelated to the service. Further, respondents indicate whether the attribution to car2go® is either exclusive or partial, leading to estimates of “narrow” and “broad” causal effects of the service.

Retrospective reporting may have reliability issues as recollection of past circumstances may be altered by subsequent experiences (Mokhtarian and Cao 2008). A reviewer suggested that the recollection of vehicle acquisition and disposal is less likely to be biased given that a vehicle is a large asset. However, recollection of circumstances and reasons behind those decisions, as well as recollection of other dimensions of travel behavior like transit usage, still have the potential for recall bias. Further, attribution of actions related to vehicle holdings and carsharing enrollment may be distorted by subjectivity and social desirability (Bertrand and Mullainathan 2001).

Firnkorn, Müller (2012) found that around 5% of the respondents attributed their decision to dispose of a car or defer purchase of a car entirely to the availability of the car2go® service. An additional 6% attributed their decision only partly to car2go®. Based on the numbers presented in the study, we estimate an average reduction of 0.05 to 0.11 vehicles per household that is attributable to carsharing. The study by Martin et al. (2010) found that members reduced their household vehicle holdings from an average of 0.47 to 0.24 vehicles or an average change of 0.23 units per household after enrolling in carsharing. Disaggregating the sample between U.S. and Canada carsharing members, the average changes in vehicle holdings are 0.26 and 0.18 respectively. The study attributes the total change in average holdings to carsharing and does not account for other factors (such as changes in income or household size) that could have contributed to changes in vehicle holdings. The study by Martin et al. also found a decrease in transit usage but an increase in walking, biking and carpooling.

Unlike the studies by Firnkorn, Müller (2012) and (Martin et al. 2010) which relied on member-only surveys with retrospective reporting, a few studies infer causal effects by comparing the travel behavior of members with that of non-members (control). Sioui et al. (2013) compared the travel behavior of members of the Montreal-based Communauto® service with the larger Montreal population. The study was facilitated by a purpose-built survey of members and an independent large-scale household travel survey conducted around the same time. Similarly, Kopp et al. (2015) surveyed a control group as well as members of the DriveNow® service in Munich and Berlin, and electronically recorded their travel activity over a seven-day period.. Both studies find large differences in travel behavior between members and non-members. However, these differences may not be entirely attributed to carsharing because the authors do not explicitly control for the biases discussed in Table 1.

Unlike previous studies which were based on one-off surveys, the study by Cervero et al. (2007) involves repeated cross-section or pseudo-panel surveys of members of City CarShare®, in the SF Bay Area. To draw causal inferences, the authors compare members' travel behavior with that of respondents who applied to join the program, but were not members for various reasons, including service unavailability in their neighborhoods (control group). The report suggests that "…these non-members were ideal controls because they displayed comparable levels of motivation.." (p. 14). We interpret this as a reference to controlling for differences in unmeasured confounders such as attitudes, which could potentially lead to selection bias due to unobservables. A non-parametric difference-in-differences analysis of mean changes in travel behavior before and after the service introduction allows for further control of observed and unobserved differences between members and non-members, and also controls for external influencers of travel behavior like gasoline prices or economic performance. Based on one survey before the launch of City CarShare® in 2001 and four surveys between 2001 and 2006 after the service launch, the authors seek to differentiate between short-term and long-term effects of carsharing. Cervero et al. (2007) conclude that carsharing leads to lower VMT, partly due to vehicle disposal.

The present study takes a different approach to inferring effects of carsharing. While purpose-built member-specific surveys are likely to collect detailed information about respondents' carsharing enrollment decisions and usage patterns, and address the potential for some biases through careful questionnaire design, they are proprietary and expensive. By contrast, various general purpose and publicly available regional travel surveys have started to capture some information about respondents' carsharing enrollment status and/or usage patterns, providing opportunity for a wider examination of the effects of emerging mobility options. Examples include the 2010-12 CHTS used in this paper and the annual United Kingdom (UK) National Travel Survey (2013 onwards). These surveys yield samples that are often larger and more representative of the general population, and they are likely to collect information about broader travel behavior patterns. On the downside, these surveys do not capture the dynamic aspects of travel behavior. Specifically, it is not possible to confirm the temporal sequence of events, leading to the potential for reverse causality. This, in addition to the potential for self-selection bias, necessitates application of various econometric and statistical tools to make causal inferences. We elaborate on many of these shortcomings in Section 7, and conclude that our effect estimates should be treated with caution.

3. Causal inferences

A question about the effect of carsharing may be phrased in three alternative ways (with different answers): (a) What is the likely effect of carsharing on a randomly chosen individual from the population of interest? (b) What has been the average effect of carsharing on those who have enrolled? (c) What would be the average effect if non-members enrolled in carsharing?

In the causal inference literature, the three estimands are referred to as ATE or average treatment effect; ATT or average treatment effect for the treated; and ATC or average treatment effect on the controls, respectively. In the potential outcomes framework, the three estimands may be represented as:

$$\begin{aligned} \tau_{ATE} &= E[y_i(1) - y_i(0)|x_i] \\ \tau_{ATT} &= E[y_i(1) - y_i(0)\,|x_i, d_i = 1] \\ \tau_{ATC} &= E[y_i(1) - y_i(0)\,|x_i, d_i = 0] \end{aligned} \quad \textbf{(1)}$$

where $y_i(1)$ denotes the outcome that will be realized by individual i if she participates in the program (enrolls in carsharing). Correspondingly, $y_i(0)$ will be realized if she does not participate (does not join carsharing). Individual i may or may not enroll in carsharing. If she enrolls (*d*=1), then $y_i(1)$ will be realized and $y_i(0)$ will ex-post be a counterfactual outcome. Similarly, if she does not enroll (*d*=0), $y_i(0)$ will be realized and $y_i(1)$ will be the counterfactual outcome. Given the above, τ_{ATT} and τ_{ATC} represent the average treatment effect over the subpopulation of carsharing members and non-members respectively.

In this paper, we focus on the second estimand – the effect on existing members or ATT. We discuss this further in Section 5.

4. Data

This study utilizes data from the 2010-2012 California Household Travel Survey (CHTS), which collected detailed information about household demographics and travel activity for the purposes of modeling statewide and regional travel and carbon emissions (Caltrans 2013). Respondents were asked whether they were members of a carsharing organization. Around 800 of the 84,000 individuals 18 years or older indicated they were members, and 80% of these carshare members were employed. Given the strong association between employment status and travel behavior, and only a small sample of carsharing members who were not employed, we restricted our analysis to employed respondents. Further, we only considered employed respondents who had identified a fixed primary work location outside their homes. This study focuses on the San Francisco Bay Area, which contains 50% of the respondents statewide who indicated being enrolled in carsharing. Further, for the sake of simplicity, we dropped cases with incomplete data on travel behavior outcome variables and/or the key covariates, thereby implicitly assuming missingness completely at random. The final sample, prior to matching, consisted of 8,299 respondents including 241 carsharing members (see Mishra et al. (2015) for more information).

The travel behavior variables modeled in this study include vehicle holdings, transit usage, and frequency of bike and walk (non-motorized) trips (all but the first measured as number of trips in a week). Number of household vehicles (y_V, HHVeh) is modeled as an ordinal variable with four levels – 0, 1, 2, and 3 or more vehicles per household. Frequency of transit (y_T) and non-motorized trips (y_N) are modeled as continuous variables after a logarithmic transformation. There are multiple alternatives for representing the travel behavior indicators. For example, vehicle ownership may be treated as continuous or ratio-scaled (which would have allowed, say, 1.34 vehicles in a household), as ordinal, as nominal (unordered categorical), or as count variables. All of these approaches have been used to model integer-valued travel behavior indicators. In the context of non-recursive SEM, however, it is uncommon to have multiple count or nominal endogenous variables, or even multiple *ordered* endogenous variables, given issues around tractability and convergence. To ensure tractability, we treated vehicle ownership as an ordered variable, and transit and non-motor usage as continuous variables. The ordinal variable (compared to using a count model) also allows for the fact that there are only a finite number of categories (i.e. that we combine four or more vehicles into a single category), rather than the countably infinite set of non-negative integers.

Given the discrete ordinal nature of the vehicle holdings variable, the parametric analysis (Section 5) models the underlying latent variable u^*. For example, u_V^* may be considered as a measure of the utility of owning a certain number of vehicles, which rises with number of vehicles owned, and is modeled as a function of all the respondent characteristics listed in Table 2 (Greene and Hensher 2010; Bhat and Pulugurta 1998). We represent the carsharing membership indicator as y_C (and sometimes alternatively as d as is common in the causal inference literature), and the corresponding continuous latent propensity to join carsharing as u_C^*. Together, the ys and u^*s are called endogenous variables in the SEM parlance.

5. Methods

5.1 Matched sampling

We use propensity score-based matching to identify a sample whose distribution of covariates closely matches that of the treatment group. In this method, one or more control units are selected for each treatment unit such that the units have the same or similar values for various observed covariates. In other words, through targeted trimming of the sample, the distribution of covariates in the control group is reshaped to closely match the distribution of covariates in the treatment group, thereby minimizing any selection bias arising due to observed covariates.

To create a matched sample, propensity scores for each individual were calculated using a probit model based on various observed covariates summarized in Table 2. We then use a nearest-neighbor matching algorithm which finds 20 or fewer control matches for each individual in the treatment group based on the smallest propensity score differences. Subjects (both treated and control) without a close enough match in the other group are discarded. We set a caliper width of 0.2 while undertaking the nearest-neighbor matching, i.e. the maximum difference in propensity scores between members of a matched pair is 0.2 standard deviations of the score.

The matching process was undertaken with replacement of the control units, implying that each control response is matched to a variable number of carsharing member responses. To account for the frequency with which each control unit was matched, weights are calculated such that each matched control unit has weight proportional to the number of treatment units to which it was matched, and the sum of the control weights is equal to the total number of matched control units (Ho et al. 2013). These weights are used in subsequent estimation of summary statistics and parametric modeling.

The process of matching simultaneously removes observations outside the common support (thereby addressing bias B1) and (approximately) balances the observed covariates (thereby addressing bias B2). The probability mass distributions of covariates in the entire matched sample matches the distribution of covariates in the treatment group; as a result the coefficients estimated in the second-stage parametric model (Section 5.2) may be interpreted as

ATT (Angrist and Pischke 2008, Section 3.3). The details of matching are discussed in (Mishra et al. 2015).

Table 2. Key variables used in this study

Variable	Description
Outcomes $^{+}$	
# of HH vehicles (y_V)	Number of vehicles owned / leased by household (HHVeh). Ordinal variable with four categories. Modeled as a mediator to measure impact on other outcomes listed below.
Transit trips per week (y_T)	Number of one-way transit trips made by respondent in the week prior to survey date. Natural log-transformed (after adding 1 to prevent taking the log of 0) to achieve a distribution closer to normal.
Non-motorized trips per week (y_N)	Number of times a person walked or biked outside, including trips taken for exercise, in the week prior to survey date. Natural log-transformed (after adding 1).
Treatment Indicator	
(y_C or d)	Indicator of carsharing membership (treatment).
Covariates (Household level)	
Child	Number of children (16 years or younger) in HH.
HH employees	Number of HH members reporting to be workers.
HH bicycles	Number of bicycles in the household.
Income^{+}	Household (HH) income: 5 categories ranging from \$50,000 or less to \$250,000 or more. The relevant survey question (Caltrans 2013, Appendix 1.0, p. 40) does not specify whether the income is 'gross' or 'net'.
Own	Home owned (1) or rented (0).
Residence type	Residence type: 3 categories (single family, multi-unit housing, others including mobile home and boats).
Residential neighborhood accessibility Work neighborhood accessibility	Composite indices computed by www.walkscore.com™ reflecting the accessibility of five categories of destinations for the residential / workplace census tracts: educational (e.g., schools), retail (e.g., grocery, drug, and convenience), food (e.g., restaurants), recreational (e.g., parks) and entertainment (e.g., theaters). Ranges from 0 to 100 (high accessibility). See Mishra et al. (2015) for more details.
Residential neighborhood density Work neighborhood density	Composite indices reflecting residential and population density of the census tracts in which subject resides / works respectively. The indicator was estimated using exploratory factor analysis of five manifest variables from U.S. Census American Community Survey (ACS) 5-year estimates. These variables are (i) Percentage of housing structures in the census tract that are one-unit houses, (ii) Median number of rooms per house, (iii) Population per square mile, (iv) Percent of households in the tract with greater than 3 vehicles, (v) Percent employed who take public transportation to work. Ranges from -2 (low density) to +4 (high density). See Mishra et al. (2015) for more details.
Carsharing station in neighborhood?	An indicator of whether a carsharing station is located in the residential census tract (1=yes; 0=no).

Covariates (Individual level)	
Age	Age in years.
Commute distance	Commute distance in miles. Gives the driving distance from centroid of residential census tract to that of work location tract and calculated using MapQuest. If residential and work census tracts are the same, we assume a distance of 2.5 miles. The distribution is lognormal; hence we consider the natural logarithm.
Disability	Indicator of temporary or permanent physical condition or disability that makes travel outside home difficult (1=yes; 0=no).
Education^{+}	Education level reflecting highest degree or level completed. 4 categories from "High school or less" to "Graduate degree".
Male	Gender: male (1) or female (0).
License	Indicator for a valid driver's license (DL) (1=yes; 0=no).
Occupation	Occupation, aggregated to 3 categories from around 25 in CHTS: white-collar (CHTS "occup" code 11-31), blue-collar (32-53), and military & others (55 & 97).
# days commuted	Days in a week person commutes to work location (1 to 7).

Notes: (+) Modeled as ordinal variables.

Outcome variables (y_V, y_T, and y_N) are not included in the propensity score model (Ho et al. 2007). Further, variables that are influenced by the outcome variables of interest – specifically commute mode choice, which is affected by vehicle holdings – are excluded while forming matched pairs because their inclusion would lead to bad conditioning and induce bias B5.

5.2 Parametric model

To model the two-way relationship between carsharing and travel behavior, we adopt a simultaneous equation / path analytic model which is summarized in Equation (2) using the LISREL notation (Bollen and Noble 2011; Bollen and Pearl 2013), and is illustrated in Figure (1) using the RAM or reticular action model notation:

$$\boldsymbol{u}_{\boldsymbol{i}}^{*}(\text{or } y_i) = \boldsymbol{\alpha} + \boldsymbol{B}\boldsymbol{u}_{\boldsymbol{i}}^{*}(\text{or } y_i) + \boldsymbol{\Gamma}\mathbf{x}_{\mathbf{i}} + \boldsymbol{\zeta}_{\boldsymbol{i}}. \qquad \textbf{(2)}$$

Eq. (2) succinctly summarizes four equations with u_V^*, u_C^*, y_T, and y_N as the dependent variables, hence we have noted the endogenous variables as $u_i^*(\text{or } y_i)$**.** $\mathbf{x}_{\mathbf{i}}$ is a vector of exogenous variables associated with individual i. **B** is the matrix of coefficients that gives the expected effects of $\boldsymbol{u}_{\boldsymbol{i}}^{*}(\textbf{or } \boldsymbol{y}_i)$ on $\boldsymbol{u}_{\boldsymbol{i}}^{*}(\textbf{or } \boldsymbol{y}_i)$. The latent propensities $\boldsymbol{u}_{\boldsymbol{i}}^{*}$ are assumed to be continuous and normally distributed. The diagonal elements of matrix **B** are zero, thereby excluding a variable from having a direct effect on itself. A number of off-diagonal elements may be assumed to be zero to indicate one-way relationships between certain endogenous variables; $\boldsymbol{\alpha}$ is the vector of intercepts, and $\boldsymbol{\Gamma}$ is the matrix of coefficients that gives the expected impacts of exogenous variables on the endogenous variables. $\boldsymbol{\zeta_i}$ is the vector of equation disturbances, which captures the net effects of all unobserved influences on the **u***s (or **y**s). By design, the mean of the disturbances is zero $[E(\boldsymbol{\zeta_i}) = \mathbf{0}]$, and it is assumed that the disturbances are uncorrelated with the exogenous variables $[COV(\mathbf{x_i}, \boldsymbol{\zeta_i}) = \mathbf{0}]$. The disturbances of all endogenous variables are assumed to be normally distributed (ordered probit model or linear regression depending upon the variable). The variances of the disturbances of u_V^* and u_C^* are fixed at 1 to enable identification (Xie 1989).

The dashed unidirectional arcs (arrows) between vehicle holdings and the carsharing membership indicator represent the mutual causality between vehicle holdings and carsharing membership, thereby controlling for (potential) simultaneity. The dashed bidirectional arc connecting the disturbances of u_V^* and u_C^* represents our assumption that the corresponding endogenous variables or outcomes share at least one omitted cause; in other words the correlation between these variables cannot be fully explained by the unidirectional arcs in our model and random errors (Kline 2012). In an alternative RAM notation, these omitted variables may be represented by a latent variable, and unidirectional arcs from this latent variable to the endogenous variables reflect the "aggregated" effect of these omitted variables. Either way, this formulation largely controls for omitted variable bias in the sense that the impacts of the remaining, observed, variables can be more accurately estimated; it does not, of course, compensate for the loss of explanatory power due to the omission of relevant causes. Given the pattern of disturbance covariances in Figure 1, the variance-covariance matrix of $\boldsymbol{\zeta_i}$, called $\boldsymbol{\Psi}$, has a combination of zero and non-zero off-diagonal elements.

Figure 1. Model to investigate impact of carsharing membership on travel behavior

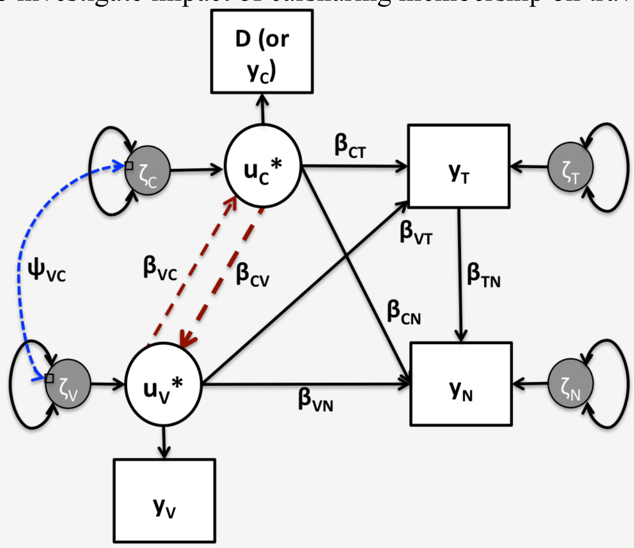

Note: Only the endogenous variables (u*s and ys) are shown; the exogenous variables (***x***) – household and individual demographics, and built environment measures for residential and job location – have been abstracted away for simplicity. Latent unobserved variables (u*s and ζs) are depicted in circles, and observed variables (ys) in squares. Subscripts C, V, T and N represent carsharing enrollment status, household vehicle holdings, frequency of transit usage, and frequency of walking and biking (non-motorized modes) respectively.

Absence of unidirectional arcs between exogenous variables and bidirectional arcs between disturbances usually reflects assumptions about the lack of a substantial relationship. For example, as seen in Figure 1, we do not model the potential influence of one's transit usage patterns on the propensity to join carsharing. Similarly, we do not account for the potential omitted causes influencing both transit usage and vehicle ownership, say personal attitudes or local policies that encourage people to substitute transit for vehicles. These omissions, labeled "strong assumptions" (because they represent the assumption of a specific value for a parameter, namely "0", as opposed to a range of possible values, such as "not 0") by Bollen, Pearl (2013), result from the high "cost" of their inclusion in the model. For example, the addition of "…each disturbance covariance to the model costs one degree of freedom and makes the model more complicated…" (Kline 2012, pp110). Similarly, unidirectional arcs in opposite directions (feedback loops) make the model non-recursive (or reciprocal) and not all non-recursive models can be identified. Stringent order and rank conditions need to be satisfied for identification of such models. In this light, the absence of various relationships in our model reflects limitations rather than freely-held assumptions.

The order and rank conditions necessary for identification of non-recursive models require the availability of exogenous covariates such that each endogenous variable in a feedback loop has a unique pattern of direct effects on it distinct from other endogenous variables (Kline 2011, Chapter 6). Given that most carsharing membership in the U.S. was neighborhood-based during the time the CHTS was conducted (Shaheen et al. 2009), an indicator of carsharing station availability in the neighborhood may be argued to influence the propensity to join carsharing but not (directly) the level of vehicle holdings (the two are likely to be correlated given that carsharing is offered in dense urban locations and vehicle holdings vary negatively with density, but there would be no reason for the availability of a carsharing station *per se* to influence vehicle holdings, except insofar as it led to carsharing membership).

Unfortunately, carsharing station availability in a neighborhood was not collected in the CHTS. We cannot establish availability based on current station locations given that carsharing operators have grown substantially since the CHTS data was collected, and historical data on station locations are not readily available. Accordingly, we estimated the indicator using the following logic: (a) if there are three or more members in a census tract, a station is assumed to exist for that tract; (b) for the cities of San Francisco, Berkeley, Palo Alto and Oakland, we relaxed the above requirement to two members or more. This analysis was done on a sample of all carsharing members, before removing respondents who were not employed or who had incomplete data on key travel behavior outcome variables. Although this is an ad hoc method for imputing station availability, we confirmed that for each tract predicted to have a station using this approach, a station does in fact exist now (although it may not have existed when the data were collected). Conversely, for all tracts (represented in the sample) that failed the test, we confirmed that a station does not exist now (and therefore probably did not exist when the data were collected, since few or no stations have closed in the Bay Area). Thus, the rules appear to be appropriate.

With respect to the other endogenous variables in the feedback loop in our model, variables that are assumed to have a direct effect on vehicle holdings but not on propensity to join carsharing are household income levels and number of children in the household. These assumptions are based on the statistically non-significant coefficients for these variables in the propensity score model for carsharing membership and correspondingly limited differences on these variables between participants and non-participants in the entire sample of employed SF Bay Area adults.

In addition to the assumed presence of omitted variables explaining some of the correlation between two of the endogenous variables (represented using the bidirectional arc in Figure 1), we also assume the presence of omitted variables that explain some of the correlation between residential location choice (specifically density index – not shown in the figure, as an exogenous variable) and both household vehicle holdings and enrollment in carsharing.

The feedback loops in Figure 1 do not reflect a hypothesis that the causal effects are truly "simultaneous", with instantaneous cycling for feedback loops. Rather, it is a limitation arising due to the cross-sectional nature of the dataset, where the variables have been measured at the same point of time. In the presence of panel data, our model would have substituted the feedback loops with a causal chain indicating that vehicle holdings in time period t-1 influence propensity to join carsharing in time t, which in turn influences vehicle holdings in time t+1. Further, vehicle holdings in t-1 and t influence holdings in t+1, reflecting habit formation or "stickiness".

The propensity u_V* may be mapped to the observed discrete levels of vehicle ownership (y_V):

$$y_{V,i} = k \qquad \text{if } \delta_{k-1} < u_i^* \leq \delta_k;\ k = 0, 1, 2, 3;\ \delta_{-1} = -\infty \;\&\; \delta_3 = +\infty\ , \tag{3}$$

where k labels the ordinal categories of ownership, and δ_k reflects the thresholds (assumed constant across all respondents) demarcating the boundaries of the continuous ranges of

u_i^* corresponding to the discrete values of $y_{V,i}$. The probability that a respondent i chooses a particular level of vehicle ownership can then be represented as:

$$\Pr(y_{v,i} = k \big| \boldsymbol{x_i}, u_{C,i}^*\big) = \Phi(\delta_k - \beta_{CV} u_{C,i}^* - \boldsymbol{\gamma x_i}) - \Phi(\delta_{k-1} - \beta_{CV} u_{C,i}^* - \boldsymbol{\gamma x_i}), \tag{4}$$

where the corresponding constant term α in eq. (2) has been absorbed into one of the threshold parameters without loss of generality.

For k = 0 (zero household vehicles), the probability may be expressed as $\Pr(y_{v,i} = 0 \big| \boldsymbol{x_i}, u_{C,i}^*\big) = \Phi(\delta_0 - \beta_{CV} u_{C,i}^* - \boldsymbol{\gamma x_i})$. Similarly for k = 3 (3 or more household vehicles), the probability may be expressed as $\Pr(y_{v,i} = 3 \big| \boldsymbol{x_i}, u_{C,i}^*\big) = 1 - \Phi(\delta_2 - \beta_{CV} u_{C,i}^* - \boldsymbol{\gamma x_i})$.

We take the (partial) effect of carsharing for an individual i on household vehicle holdings to be the difference in probabilities (eq. (4)) at the average propensity to enroll in carsharing for members $E(u_{C,i}^* | x_i, d_i = 1)$ and non-members $E(u_{C,i}^* | x_i, d_i = 0)$, both computed at $X = \boldsymbol{x_i}$:

$$\Delta\Pr(y_{v,i} = k | \boldsymbol{x_i}) = \Pr(y_{v,i} = k | \boldsymbol{x_i}, \widehat{u_{C|1}^*}) - \Pr(y_{v,i} = k | \boldsymbol{x_i}, \widehat{u_{C|0}^*}) . \tag{5}$$

where $\widehat{u_{C|1}^*}$ and $\widehat{u_{C|0}^*}$ are the estimated average propensities to enroll in carsharing for members (d=1) and non-members (d=0) respectively. The average partial effect (APE) of carsharing averages the individual effects across all individuals.

Using the estimates of potential changes in probabilities of belonging to various vehicle categories resulting from carsharing enrollment, we estimate the expected absolute change in number of vehicles owned by a household. To simplify our calculations, we assume three vehicles for the highest vehicle ownership category (K_V = 3 or more vehicles). Thus, the potential change in vehicle ownership for individual i may be estimated as:

$$\Delta\text{HHVeh}_i = \sum_{k=0}^{3} \Delta\Pr(y_{v,i} = k|\boldsymbol{x_i})\times k \tag{6}$$

or, equivalently, the difference between the expected number of vehicles individual i would own with, and without, carsharing membership. The above change is averaged across all individuals to estimate the average change in vehicle ownership due to carsharing.

This approach highlights a key limitation of estimating the causal effects of a discrete treatment using a non-recursive or reciprocal model. As shown by eq. (2), we assume that u_v^* is a function of u_c^* rather than y_C, *even though it is arguably the practical implications of membership itself, not merely the propensity to belong, which influences vehicle ownership propensity*. In fact we *must* make this assumption for the system to be reciprocal: as Maddala (1986, Section 5.7) demonstrates, the alternative assumption leads to a logical inconsistency (because inserting the equation for u_v^* into the equation for u_c^*, and simplifying, leads to an equation in which the variable y_C – which can only take on the values 0 and 1 – can be directly expressed as a function of its corresponding latent variable u_c^* – which is continuous-valued). Accordingly, to estimate the effect of enrollment in carsharing on vehicle holdings as shown in eq. (5), we must use averages of the latent continuous variable (u_C^*) representing the propensity to enroll in carsharing, and not the binary 0/1 indicator representing whether a respondent participated or not in the program.

As a result, in the context of a discrete treatment, addressing simultaneity bias with a recursive equation system apparently requires an unpleasant sacrifice of conceptual fidelity. The necessity of this compromise arises due to the combination of (1) using cross-sectional data to represent dynamic effects that occur across a period of time, and (2) the lack of random assignment of the treatment to individuals. If longitudinal data were available (so that the mutual causality was not modeled as occurring simultaneously), or if treatment assignment were random (thereby obviating the need for propensity score matching or other such approaches to mimicking random assignment), the problem would not arise.

We adopt the maximum likelihood estimator using numerical integration and cluster-robust standard errors with clustering on the residential census tracts. Residents in the same neighborhood are likely to have more similar travel behavior and other socio-economic characteristics than those of different neighborhoods. Clustering also arises because the attributes of a census tract – residential accessibility and density scores – do not vary across individuals within each tract (Angrist and Pischke 2008; Moulton 1990)

Finally, we should note that the probit model adopted for the propensity score model is not the same as the one adopted for the path analytic model in Figure 1 with carsharing membership (y_C) as the outcome variable. The propensity score model did not include vehicle holdings (y_V) given that it is an outcome variable of interest (Brookhart et al. 2006). Aside from that omission, the propensity score model is less parsimoniously specified than the path analytic model, consistent with the advice in the literature (Austin 2011; Rubin 2001) that collinearity and causal hypotheses are of little concern in specifying propensity score models.

5.3 Direct and mediated (or indirect) effects

The coefficients β_{CV}, β_{CT} and β_{CN} capture the "direct" effects of carsharing. In addition to affecting the frequency of taking transit and non-motorized trips along "direct" pathways, carsharing influences these dimensions of travel behavior along "indirect" or "mediated" pathways. For example, carsharing influences the frequency of bike and walk trips along the following three indirect pathways: (i) mediated by vehicle holdings ($u_C^* \rightarrow u_V^* \rightarrow u_N^*$), (ii) mediated by transit usage ($u_C^* \rightarrow u_T^* \rightarrow u_N^*$), and finally (iii) mediated by both vehicle holdings and transit usage ($u_C^* \rightarrow u_V^* \rightarrow u_T^* \rightarrow u_N^*$). The total effect of carsharing is the sum of the direct and the abovementioned mediated effects (Muthén 2011).

5.4 Alternative model specification

In Section 5.2 above, we estimated the parametric model on a matched sample where the distribution of covariates in the control (non-member) group is similar to that in the treatment (member) group. The estimand of interest in that section is the average treatment effect on the treated.

In an alternative model, we estimate the parametric model on the entire sample of employed San Francisco Bay Area residents (8,299 observations). As before, the coefficient β_{CV} represents the causal effect of carsharing on vehicle holdings. However, given that the probability mass functions of the covariates differ across the two treatment groups, β_{CV} may not be comparable to any of the three estimands listed in eq. (1) (Angrist and Pischke 2008; Morgan and Winship 2014). When using the entire unmatched sample, Angrist, Pischke (2008, Section 3.3) show that the regression estimate of the treatment effect may be written as

$$\beta = \frac{\sum_x \tau_x [P(D_i = 1|X_i = x)(1 - P(D_i = 1|X_i = x))] P(X_i = x)}{\sum_x [P(D_i = 1|X_i = x)\,(1 - P(D_i = 1|X_i = x))]\; P(X_i = x)} \tag{7}$$

i.e. as the variance-weighted average of τ_x over the entire distribution of **X**, where τ_x is the estimated effect at a specific value of the covariate vector **X** = x, and the variance in question is the variance of the binary treatment status variable, conditional on X = x. This approach puts the maximum weight on covariate "cells" where the conditional variance of treatment status is largest, which is when the probability of treatment is ½; i.e. for those values of covariates where there are equal numbers of control and treatment members. Although it may be logical to give the greatest weight to cells that most closely resemble a random assignment to control and treatment groups, Morgan, Winship (2014, p. 206) note that such variance-weighted average effects may not be of "…any inherent interest to a researcher …" because they cannot be related to the important estimands in the causal literature – ATT, ATE and ATC.

Because the matching process results in a similar ratio of members to non-members in each covariate cell, the conditional variance is the same over the entire distribution of **X**. Further, given that the probability mass function of the matched sample is similar to that of the members, the coefficient β_{CV} estimated using a matched sample may be interpreted as the ATT.

The alternative model is motivated by the common practice in travel behavior literature, where most parametric analyses with the objective of estimating causal effects proceed with the entire sample rather than a smaller matched sample. Parametric (regression) analysis may be used to control for potential confounders (biases B1 and B2) by explicitly including them as explanatory variables. In that sense, regression control may be seen as an alternative to matched sampling. However, recent literature has questioned the efficacy of regression analysis when the covariate distributions of program participants and non-participants are seriously imbalanced. For individual "important" covariates, a simple rule-of-thumb is that the mean for the two groups should be no more than one-quarter or at most one-half of a standard deviation apart (Imbens and Rubin 2015, pp 277). Rubin (2001) provides more formal and detailed criteria based on comparison of distributions of the entire multidimensional covariate space (X) represented by the scalar propensity score. When these conditions are not met, then "…linear regression on random samples gives wildly erratic results... , sometimes markedly overcorrecting [percentage bias reduction >> 100%], or even ... greatly increasing the original bias [percentage bias reduction << 0%]…" (p. 175).

Table 3 (column 4) indicates that there are large differences in covariate distributions between the carsharing members and non-members, and thus could lead to biased parametric estimates in case of a regression analysis without any exercise to balance the covariates. Accordingly, estimation of the alternative model, and comparison to the original one, allows us to assess the extent of this possible bias. We should note that matching samples is only one of the many ways to balance covariates; other approaches include propensity score-based weighting and sub-classification (Rubin 2001).

5.5 Software

Most of the initial data cleaning and analysis was undertaken using various packages in R. The matched sample was identified using the MatchIt package in R (Ho et al. 2013). The path analytic model was undertaken in MPLUS Version 7.4 (Muthén and Muthén 1998-2014).

6. Results

6.1 Matched sample and covariate balancing

The process of propensity score matching generated a matched sample of 2,302 non-members (control or non-participants) and 238 carsharing members (program participants) from a total sample of 8,512 Bay Area employed respondents. Observations for around 6,000 non-members were outside the convex hull and could not be matched to any carsharing member. These observations were dropped from the matched sample. Similarly, three carsharing members with high propensity scores could not be matched to any non-members and were also dropped.

The means and standardized mean differences of important confounding factors for members and non-members are summarized in Table 3. The table as well as density plots of various covariates shown in Mishra et al. (2015, Supporting Information) highlight that after matching, members and non-members are substantially balanced in their distribution of observed covariates. For example, the difference in standardized means of workplace accessibility score reduces from 0.85 to 0.02; or in non-standardized terms from 26 units to less than 1 unit. Similar improvement in covariate balance is observed in nearly all other covariates.

Table 3. Differences in covariates between carsharing members and non-members (comparison of means)

	Entire Sample			**Matched Sample**		
	Members	**Non-members**	**Standardized Mean Differences**	**Members**	**Non-members**	**Standardized Mean Differences**
Sample size	241	8,271		238	2,302	
Average propensity score	0.121	0.026	0.865	0.117	0.116	0.005
Annual household income ('000s)						
< $50	8%	14%	-0.254	8%	8%	- 0.031
$50-100	25%	29%	-0.096	25%	24%	0.026
$100-250	59%	48%	0.206	59%	59%	-0.037
>$250	8%	9%	-0.023	8%	9%	-0.003
% single family residences in residential census tract	53%	83%	-0.630	53%	52%	0.021
# of children in HH	0.47	0.56	-0.130	0.48	0.47	0.010
# of HH members employed	1.78	1.86	-0.087	1.79	1.77	0.028
# of HH members with driver's license	1.92	2.16	-0.333	1.92	1.90	0.026
Built environment (residence)						
Accessibility score	74.87	48.59	1.020	74.15	73.71	0.017
Density index	1.15	-0.11	1.024	1.13	1.14	-0.006
Built environment (workplace)						
Accessibility score	76.85	56.75	0.850	76.79	76.08	0.003
Density index	1.49	0.57	0.800	1.49	1.48	0.003
Commute distance (miles)	11.8	14.9	-0.208	12.00	12.08	-0.007
Age (years)	45.72	48.54	-0.262	45.79	46.03	-0.023
Male	61%	51%	0.202	61%	63%	-0.036
Education						
Undergraduate	41%	31%	0.193	41%	40%	0.026
Graduate	46%	32%	0.200	46%	46%	0.024

Note: The standardized mean differences are estimated by dividing the mean differences by the standard deviation in the original treatment group.

The covariate balancing in the matched sample allows us to claim that selection bias arising due to systematic differences in observed characteristics between members and non-members (biases B1 and B2) has been controlled. We should note that as a result of selecting "nearest" rather than "exactly" matched control units, we have not entirely eliminated the differences in covariates. Subsequent parametric analysis will allow us to control for the residual differences.

6.2 Model parameter estimates

Table 4 gives the estimates of key parameters and standard errors (in parentheses) for the model summarized in Figure 1 using the matched sample (I. Main Model) and entire sample (II. Alternative Model). The complete MPLUS results may be obtained from the corresponding author. The estimated coefficients of the former may be interpreted as the average treatment effect on the treated (ATT), while the latter gives variance-weighted average effects.

Estimates of β_{CV} are negative and statistically significant, indicating that carsharing incentivizes a decrease in average vehicle holdings - a finding consistent with the literature. For the ATT results, the estimated covariance of the disturbances of propensities u_V^* and u_C^* (ψ_{VC}) is negative, indicating that unobserved variables that incentivize enrollment in carsharing tend to reduce household vehicle holdings, and those that disincentivize carsharing membership tend to increase holdings. β_{CV} and ψ_{VC} have the same sign (negative), indicating that not controlling for omitted variables (B3) would lead to overestimation of causal effects.

The direct and total effects of carsharing on transit usage and walk and bike frequency are positive, but statistically significant only for the variance-weighted average effects. The indirect effects, mediated through reductions in vehicle holdings, are in the same direction as the direct effects. The increase in non-motorized trips may may arise due to vehicle disposal, and/or partly reflect trips between home and the carsharing station. Similarly, an increase in transit trips may arise as a consequence of vehicle disposal in connection with carsharing membership. We note that the results for both variables are not statistically significant for the main model.

Table 4. Two alternative estimates of the model summarized in Figure 1 (key coefficients and standard errors)

		I. Main model (matched sample)	II. Alternative model (full sample)
	Treatment estimand	Average Treatment Effects on the Treated (ATT)	Variance-weighted average effects
I	Sample size (members / non-members)	238 / 2302	241 / 8271
II	Estimated coefficients (direct effect of carsharing)		
	β_{CV} ($u_C^* \rightarrow u_V^*$)	**-0.235 (0.084)**	**-0.557 (0.083)**
	β_{CT} ($u_C^* \rightarrow y_T$)	0.043 (0.031)	**0.080 (0.031)**
	β_{CN} ($u_C^* \rightarrow y_N$)	0.058 (0.032)	**0.132 (0.020)**
III	Estimated total effect of carsharing (Direct as above + indirect/mediated)		
	Effect on transit usage ($u_C^* \rightarrow y_T$)	0.057 (0.031)	**0.089 (0.031)**
	Effect on walk and bike frequency ($u_C^* \rightarrow y_N$)	**0.064 (0.031)**	**0.135 (0.019)**
IV	Other important coefficients		
	β_{VC} ($u_V^* \rightarrow u_C^*$)	-0.005 (0.043)	-0.019 (0.024)
	β_{VT} ($u_V^* \rightarrow y_T$)	**-0.061 (0.019)**	**-0.015 (0.007)**
	β_{VN} ($u_V^* \rightarrow y_N$)	-0.008 (0.016)	-0.001 (0.008)
	β_{TN} ($y_T \rightarrow y_N$)	**0.070 (0.027)**	**0.029 (0.012)**
V	Effect of omitted variables		
	ψ_{VC} ($u_V^* \leftrightarrow u_C^*$)	**-0.334 (0.148)**	-0.068 (0.119)
	$u_V^* \leftrightarrow$ Density index (residence)	**-0.520 (0.098)**	**-0.166 (0.060)**
	$u_C^* \leftrightarrow$ Density index (residence)	**-0.428 (0.109)**	**0.231 (0.060)**
VI	Model goodness-of-fit indicators	None reported by MPLUS. See main text	
VII	Explained variances (R-squared)		
	Membership indicator (u_C^*)	0.643	0.553
	Vehicle holdings (u_V^*)	0.791	0.829
	Weekly transit trips (y_T)	0.428	0.305
	Bike & walk trips (y_N)	0.157	0.133

Note: Parameter estimates in bold are statistically significant at the α = 0.05 level.

Unfortunately, Mplus was not able to estimate the degrees of freedom and hence calculate the various goodness-of-fit statistics for our model given the model complexity (personal communication with MPLUS support team, December 2015). Section VII of Table 4 gives the "estimated" R-squared values for the underlying continuous latent variables of y_Y and y_C (u_C^* and u_V^* respectively), and the "actual" R-squared values for the continuous endogenous variables (y_T and y_N) (McKelvey and Zavoina 1975). The estimated R-squareds, similar to the actual R-squareds for continuous variables, indicate the extent to which the model explains the variance of the underlying latent variables, but unlike the latter are estimates rather than actual values of these quantities. The estimated R-squares for the u_C^* and u_V^* equations are quite high; none of the authors have seen such high R-squared values in the travel behavior literature pertaining to disaggregate models on cross-sectional data, but we do not have any further explanation.

6.3 Effect of carsharing

Using eq. (4), Figure 2(A) gives the average estimated probability of belonging to a specific category of HH vehicle holdings after controlling for various biases listed in Table 1. Probabilities were estimated for all respondents with and without the assumption of enrollment in carsharing, using the average propensities to enroll for the members and non-members in the sample respectively.

The probabilities are disaggregated by neighborhood type and "potential" enrollment status. We divided the sample into three groups based on the residential accessibility scores – 75th percentile and higher, median to less than 75th percentile, and less than median. The underlying accessibility scores map closely to www.walkscore.com's© definition of "Very Walkable" or higher, "Moderately Walkable", and "Car Dependent". The first two groups may be considered urban regions while the last one suburban. In each of the neighborhoods, there is an increase in the probability of not owning any vehicles or owning only one vehicle in the household as a result of enrolling in carsharing. For urban neighborhoods, there is a corresponding decrease in the probability of owning more than 1 vehicle. For suburban neighborhoods, there is a small increase in the probability of owning 2 vehicles in the household but a larger decrease in the probability of owning 3 or more vehicles.

Figure 2. Effect of carsharing on household vehicle holdings disaggregated by neighborhood type

(A) Average probabilities of various levels of HH vehicle ownership, disaggregated by potential enrollment status and neighborhood type	(B) Change in household vehicle holdings resulting from enrollment in carsharing

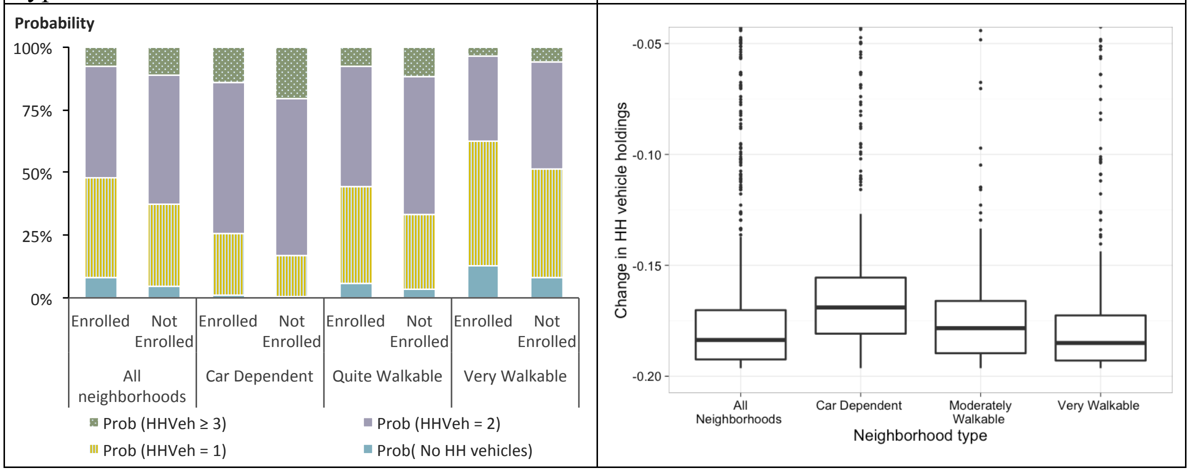

Note: The above effects represent the average treatment effect on the treated (ATT) estimated based on parametric analysis of the matched sample. Effects are disaggregated by neighborhood type – "Car Dependent" may be considered a suburban neighborhood, while "Moderately Walkable" and "Very Walkable" may be considered urban neighborhoods. See text for more details.

Figure 2(B) indicates that carsharing enrollment has led to an average fall in household vehicle holdings of 0.17 units. This is similar to estimates by both Martin et al. (2010) and Firnkorn, Müller (2012) who estimated reductions of 0.18-0.26 and 0.05-0.11 vehicles per household respectively. However, we should caution the readers about the comparability of the numbers, given large differences in the underlying datasets. Our sample is restricted to employed respondents, while the other studies included members who are not in the labor force. Further, there are differences in time period, economic conditions during the study period, geographic coverage, and most importantly, the underlying characteristics of the carsharing services. At the time the CHTS survey as well as the survey by (Martin et al. 2010) were conducted, the carsharing industry was largely a station-based service where vehicles are picked up and returned at the same location. On the other hand, the car2go® service considered by Firnkorn, Müller

(2012) is a free-floating carsharing service that provides more flexibility to members in terms of vehicle pick-up and drop-off. Further, the car2go® service offers electric vehicles while the station-based carsharing services in 2011-12 offered almost exclusively conventional vehicles. Given the difference in product features in terms of flexibility and vehicle technologies, we had initially anticipated lower effect estimates compared to estimates by Firnkorn, Müller (2012). Differences in share of respondents in labor force, as well as in the underlying rate of vehicle ownership in the two regions may potentially explain why our estimates of the impact of a station-based carsharing service are higher than those estimated for a free-floating service; however, more data and analysis are required for a more definitive conclusion.

Analysis of the parametric estimates of the alternative model specification leads to a variance-weighted causal effect of around 0.44 units; in other words nearly one in every two households that enrolled in carsharing discarded a car as a result of incentives provided by carsharing. We feel this estimate is quite high, casting doubt on the credibility of the model (also see the discussion at end of Section 5.4). We do not discuss this particular result any further.

6.3 Bias Reduction

Figure 3 highlights the role of various biases in explaining the differences in vehicle holdings between carsharing members and non-members. Panel (A) gives the average probabilities of falling into a particular vehicle holding category either observed in the sample (for the first two pairs of bars) or estimated using the methods listed in Section 5.2 (for the last two pairs). Panel (B) gives the corresponding difference in vehicle holdings estimated using eq. (6).

The first pair of columns in panel (A) gives the observed share of respondents in various categories of vehicle holdings for the entire sample of 8,512 San Francisco Bay Area employed respondents. Most non-members have 2 or more vehicles per household. On the other hand, nearly one-third of the members do not have any vehicles. The average household vehicle holdings are 1.1 and 2.0 for members and non-members respectively. The difference in vehicle holdings of 0.9 between households with and without carsharing members may be partly explained by observed differences in household and individual demographics and the built environment of the residential neighborhood (Table 2 in Section 6.1, biases B1 and B2). The 32% of carsharing members in our dataset that are from households without any vehicles may be compared to around 48% of car2go® members in Ulm (Firnkorn and Müller 2012) and 88% of Communauto® members in Montreal (Sioui et al. 2013) without any vehicles in the household. Correspondingly, Martin et al. (2010) found that in Canada and the U.S. around 62% and 80% respondents were from zero-vehicle households before and after joining carsharing, respectively. In addition to sampling variance, the difference may be potentially attributed to differences in the respective populations – our study only includes employed respondents while the other studies include members not in the labor force. The latter are less likely to own vehicles. Further, the Martin et al. study excludes "inactive" members whose lack of carsharing usage could potentially be due to access to household vehicles.

In the second pair of columns, we control for these systematic differences in observed characteristics using propensity score matching. The columns show the observed share of respondents in various categories for the matched sample of 2,540 respondents. The trimmed control (non-member) sample has an average household vehicle holding of less than 1.8 units. Since the set of members in the matched sample is largely the same as in the original untrimmed sample, there is no change in the average holding of 1.1 for members. The remaining difference of 0.7 units may be explained partly by the causal effects of carsharing, and partly by systematic differences in *unobserved* characteristics between members and non-members (omitted variable bias, B3); as well as reverse causality or simultaneity bias (B4). The fall in difference of around 0.2 units (0.9 – 0.7, green bar in panel (B)) may be attributed to the systematic differences in the distribution of *observed* characteristics between members and non-members.

Figure 3. Decomposition of the observed differences in vehicle holdings between carsharing members and non-members due to the causal effects of carsharing

(A) Observed proportions or estimated average probabilities of household vehicle holdings

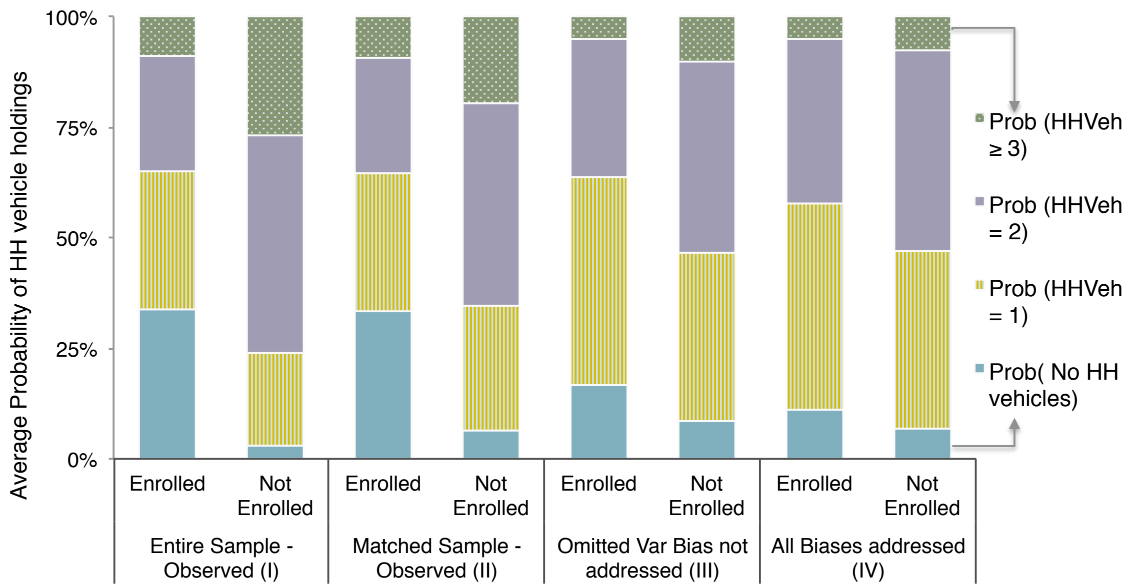

(B) Disaggregation of observed differences in vehicle holdings into causal effects and various biases

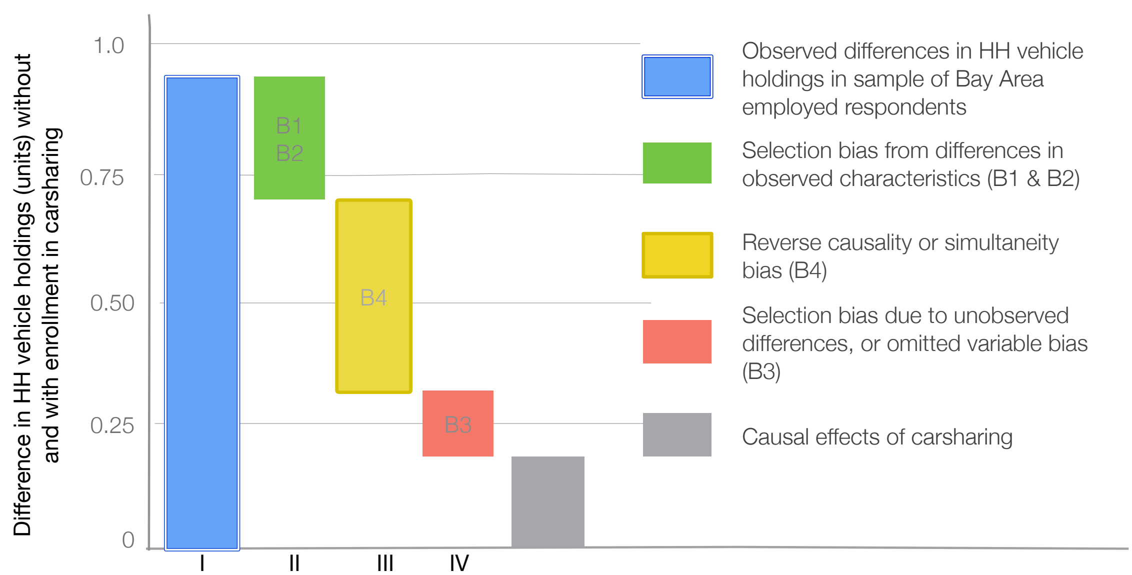

Note: The distribution of vehicle holdings in (I) and (II) in (3A) is based on observed data in the entire sample and matched sample respectively, and not based on the results of our parametric model estimates.

The third pair of columns in panel (A) reflects the probabilities of vehicle ownership with and without enrollment in carsharing, calculated based on the estimated parameters of a model very similar to Figure 1 with one key difference: ψ_{VC} is assumed to be zero. That is, the covariance between vehicle holdings and carsharing membership is assumed to be explained entirely by the unidirectional arcs in our model and random errors. As discussed before, this is a very strong assumption given that we do not have information about respondents' attitudes and preferences on such issues as the environment, driving, walking and biking – preferences which have been shown in the literature to substantially explain people's travel and residential location decisions. However, the model does control for simultaneity bias (B4) as well as selection bias due to observed differences in characteristics (B1 and B2). The estimated average household vehicle holdings after controlling for B1, B2 and B4 are 1.2 and 1.5 with and without enrollment in carsharing. The yellow bar in panel B (0.4 vehicle units) may be attributed to simultaneity bias.

Finally, the last set of columns gives the probabilities calculated based on the parametric estimates from our final model. The estimated average household vehicle holdings are 1.3 and 1.5 with and without enrollment in carsharing, where the difference of 0.17 units may be attributed to the effects of carsharing (grey bar in panel (B)).

7. Discussion and Conclusions

This study adopts a framework to draw causal inferences from cross-sectional observed data by combining well-established methods in the statistical literature to address the joint occurrence of multiple biases. In the first step, a propensity score-based matched sample is identified to control for selection bias arising from observed differences in characteristics between program members and non-members (biases B1 & B2 in Table 1). Based on the matched sample, a non-recursive structural equation model is estimated to address bi-directional effects or reverse causality, i.e. to account for the influence of the outcome itself on program participation (B4), distinct from the influence of the program on the outcome of interest. In this non-recursive model, the propensity to enroll in the program as a function of outcome and other covariates is simultaneously estimated along with the outcome as a function of the program participation and other covariates. The errors of two equations are allowed to be correlated to address omitted variable bias or self-selection bias due to unobserved influencers of both program participation and vehicle holdings (B3).

We apply the framework to estimate the effect of carsharing on the vehicle holdings and travel behavior of existing members residing in the San Francisco Bay Area using the California Household Travel Survey (CHTS) dataset. The application presents a good example in which all the biases listed in Table 1 have to be addressed. There are large differences in observed individual and household demographics and residential and workplace location choices between members and non-members, leading to the potential for selection bias due to observed differences in characteristics. Reverse causality arises because household vehicle holdings have been identified in the literature to influence enrollment in carsharing programs, and the cross-sectional nature of CHTS fails to capture this dynamic process. Further, the potential for omitted variable bias or selection bias due to unobservables arises because CHTS does not measure respondent attitudes which are believed to influence travel behavior outcomes of interest and carsharing enrollment.

We find that 80% of the observed difference of 0.9 units in average vehicle holdings between carsharing non-members and members may be explained by the biases listed above. The remaining difference of 0.17 units reflects the estimated effect of carsharing, which is the equivalent of shedding one vehicle by about one out of every six households whose member(s) are enrolled in carsharing. The effects are larger in urban neighborhoods, where members are likely to have easier access to the neighborhood-based carsharing stations compared to suburban members. Carsharing also modestly incentivizes increased frequency of biking, walking and transit use, both directly, and indirectly mediated by reductions in vehicle holdings.

The methodological framework demonstrated in this study may be adopted to draw causal inferences from other programs with bi-directional effects, using cross-sectional observational data. Such programs in the travel behavior literature include telecommuting and its effect on (say) residential location choice and one-way commute distance. Commuters with long commute distances will prefer to telecommute; on the other hand, telecommuters have an incentive (or lower costs) to choose residential locations farther away from their workplace. However, the framework has several limitations. We discuss these limitations in the context of carsharing, but the discussion is relevant for other applications also.

The first limitation arises as a result of the particularities of the travel behavior problem being studied, where vehicle holdings both influence the propensity to enroll in carsharing, and are influenced by the membership status of a respondent. Specifically, we use a cross-sectional dataset to estimate effects of a dynamic process that plays out over a period of time, as discussed earlier in Section 5. This necessitates a non-recursive path model with feedback loops, which in turn requires the key assumption of equilibrium, meaning that any causal impacts (travel behavioral changes) have manifested and the system is in a steady state. Violation of the equilibrium assumption can lead to severely biased results (Kline 2011, pp108). From the available dataset, it is not feasible to assess the stability of the system. It is likely that the membership duration varies across respondents – in which case causal effects are at different stages of manifestation and hence may not have reached equilibrium.

In addition to the strong assumption of equilibrium, identification of treatment effects in a non-recursive model requires the availability of exogenous covariate(s) that influence(s) the propensity to enroll in carsharing but not vehicle holdings. As discussed in Section 5.2, we depended upon a variable that was synthetically estimated by us based on the sample distribution and not originally collected as part of CHTS – an indicator of carsharing station availability in the respondent's residential neighborhood. Given that most carsharing membership was neighborhood-based during the time the CHTS was conducted, an indicator of carsharing station availability in the neighborhood may be argued to influence the propensity to join carsharing but not the level of vehicle holdings. Although we adopted very strict criteria to estimate the indicator, the synthetic nature of the variable gives us reason for caution in adopting our estimates of causal effects.

The above discussion highlights the need for caution while interpreting causal estimates from non-recursive models. Per Bagozzi, Yi (2012 pp31) "…A problem with reciprocal causation models is that it is difficult to justify mutual causality between [two endogenous variables] when these are measured at the same point in time. At least to the extent that temporal priority of cause to effect is a defining quality of causation, such a model may not be defensible, although it can be estimated and tested with SEM procedures. Causality is regarded fundamentally to be a recursive phenomenon under most conceptualizations in philosophy…"

The study assumes homogeneity in treatment effects – specifically, the effect coefficients β_{CV}, β_{CT}, and β_{CN} are assumed to be invariant across respondents. In an alternate, and definitely more plausible, scenario where treatment effects vary across individuals, then the assumption of homogeneous effects may lead to biased estimates (Xie 2013).

Yet another limitation of a non-recursive model (as discussed at greater length in Section 5.2) is that, because the treatment (membership) indicator is endogenous, it cannot take the intuitive binary 0/1 form when it appears as an explanatory variable in the equation for vehicle holdings (eq. (2) and Figure 1). To estimate the effect of enrollment in carsharing on vehicle holdings (eqs. (5) and (6)), we must use the means of the latent continuous variable (u_C^*) representing the *propensity* to enroll in carsharing. This approach is necessitated by the departure from the paradigm of true experiments, where treatment is assigned randomly and is independent of the subjects' inclination. Various econometric methods that seek to measure causality from observational data, including propensity score matching, aim to mimic this method of random assignment. However, the necessity of replacing the observed treatment indicator with its estimated latent propensity is at odds with the logic that it is the treatment itself (not the propensity to be treated) that is the true cause. Hence the use of latent propensities to enroll in carsharing as an endogenous explanatory variable seems flawed and is a compromise.

In future, researchers can adopt a continuous treatment variable (if available, which was unfortunately not the case for us) to address the above shortcoming. For example, instead of the member/not-a-member indicator adopted in this study, future travel surveys could collect data (say) about number of carsharing trips made in a week or month to capture both enrollment (treatment status) and usage (level of treatment). In a medical context, a similar example is using level of drug dosage instead of an indicator that only captures whether the drug was given or not. However, to the extent that the estimated probability of treatment is correlated with the "intensity" or frequency of treatment, using the continuous-valued propensity variable in lieu of the 0/1 treatment indicator is not necessarily bad (and in fact may be better than using the cruder 0/1 variable). Such a correlation would undoubtedly not be perfect or nearly so, but is likely to be substantial. We should also note that only a small percentage of papers in the causal inferences literature have used a continuous treatment variable, partly because of lack of data availability, and partly because of difficulties in undertaking propensity score matching when multiple treatment levels are present.

The study by Martin et al. (2010) suggests heterogeneity in circumstances influencing enrollment decisions; this is distinct from the heterogeneity in treatment effects discussed earlier. Some members may have joined carsharing to enhance their mobility options. Martin et al. (2010) found that around 50% of members in North America did not own a vehicle and joined carsharing to get access to vehicles. We speculate that this group of members may be expanded to include suburban residents without carsharing in their neighborhoods, but who use carsharing at their workplace and when visiting city centers. For this group, it may be hypothesized that carsharing influences increased driving and trip making, but not a reduction in vehicle holdings. A second group of members joins carsharing to facilitate reduction in vehicle holdings. For example, in the same study by Martin et al., around 30% indicated that they joined carsharing to avoid purchasing an additional vehicle or to dispose of their vehicle. For this group, it is evident that carsharing influences, or at least facilitates, a reduction in vehicle holdings. Future studies should seek to address such heterogeneity. The method adopted by Waddell et al. (2007) suggests one promising alternative, whereby the sample is probabilistically partitioned into two segments using latent class methods, and a recursive model is fitted to each segment. The two recursive models specify causality in opposite directions.

The above discussion highlights various threats to the internal validity of our results. In addition, we should also caution readers about the external validity of our results. We initially restricted our population of interest to employed respondents from the San Francisco Bay Area. The process of matching effectively narrowed our population even more, to those having the specific characteristics that were summarized in the last three columns of Table 3. This population largely resides and works in urban neighborhoods, is predominantly male and highly educated, and has relatively short commute distances. Even further, our effect estimates may be interpreted as the average treatment effect on the *treated* (ATT) members of that specific subpopulation, meaning that even within that subpopulation, we cannot necessarily expect the same effects for non-members should they eventually adopt carsharing. Thus, the estimated effects may certainly not be extrapolated to the entire population of Bay Area employed residents, and some idea of how modest the aggregate effects might be when diluted by the entire population is conveyed by the fact that even among the “most-likely-to-adopt” subpopulation described above, only about 12% actually become carsharing members (Table 3, average propensity score).

External validity may also be limited due to changes within the carsharing and related shared mobility services since the time the CHTS survey was conducted in 2011-12. As mentioned earlier, the industry has evolved from being largely a station-based service with only conventional vehicles to one that offers a wide range of product features including more flexibile one-way and free-floating services, and often electric-drive vehicles. Further, on-demand or ridesharing services (e.g. Uber and Lyft), which provide alternative mobility options to users, have been introduced. The ridesharing industry has grown much faster in terms of users and already has substantially more users than the carsharing industry (Clewlow 2016). It is not clear how the two alternative mobility options of carsharing and ridesharing interact with each other – however, we believe that the effects of carsharing on travel behavior may potentially be affected by the availability of ridesharing services.

The discussion of threats to the internal validity of our results highlights the difficulties in making causal inferences regarding programs with bi-directional effects, using cross-sectional observational data. However, randomized experiments or collection of long-term longitudinal data are quite rare due to cost and feasibility considerations. In their absence, the methods used in this study may provide a ballpark estimate of causal effects, provided analysts and readers are well aware of the limitations.

8. Acknowledgements

This study has benefited from the input and critical comments of several anonymous referees, and from various researchers including Kenneth Laberteaux from the Toyota Research Institute of North America; and Susan Handy, Gil Tal, Sonia Yeh, Lewis Fulton, David Bunch and Joan Ogden, all from the Institute of Transportation Studies at UC Davis; and James Sweeney from Stanford University. We are also grateful to Karim Hamza for help with obtaining accessibility scores. The views and opinions expressed in this paper are those of the authors alone. The authors declare no conflict of interest.

9. References

Angrist, J., Pischke, J.-S.: Mostly harmless econometrics: An empiricist's companion. Princeton University Press, (2008)

Austin, P.C.: An introduction to propensity score methods for reducing the effects of confounding in observational studies. Multivariate Behavioral Research **46**(3), 399-424 (2011). doi:DOI:10.1080/00273171.2011.568786

Bagozzi, R.P., Yi, Y.: Specification, evaluation, and interpretation of structural equation models. Journal of the Academy of Marketing Science **40**(1), 8-34 (2012). doi:DOI: 10.1007/s11747-011-0278-x

Bertrand, M., Mullainathan, S.: Do people mean what they say? Implications for subjective survey data. The American Economic Review **91**(2), 67-72 (2001)

Bhat, C.R., Pulugurta, V.: A Comparison of Two Alternative Behavioral Choice Mechanisms For Household Auto Ownership Decisions. Transportation Research Part B: Methodological **32**(1), 61-75 (1998)

Bollen, K.A., Noble, M.D.: Structural equation models and the quantification of behavior. Proceedings of the National Academy of Sciences **108**(Supplement 3), 15639-15646 (2011). doi:DOI:10.1073/pnas.1010661108

Bollen, K.A., Pearl, J.: Eight Myths about Causality and Structural Equation Models. In: Morgan, S.L. (ed.) Handbook of Causal Analysis for Social Research. pp. 301-328. Springer, (2013)

Brookhart, M.A., Schneeweiss, S., Rothman, K.J., Glynn, R.J., Avorn, J., Stürmer, T.: Variable Selection for Propensity Score Models. American Journal of Epidemiology **163**(12), 1149-1156 (2006)

Caltrans: 2010-2012 California Household Travel Survey Final Report. In. California Department of Transportation, (2013)

Cervero, R., Golub, A., Nee, B.: City CarShare: Longer-Term Travel Demand and Car Ownership Impacts. Transportation Research Record: Journal of the Transportation Research Board **1992**(1), 70-80 (2007)

Clewlow, R.R.: Carsharing and sustainable travel behavior: Results from the San Francisco Bay Area. Transport Policy **Advance online publication** (2016). doi:doi:10.1016/j.tranpol.2016.01.013s

Elwert, F., Winship, C.: Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable. Annual Review of Sociology **40**, 31-53 (2014). doi:DOI: 10.1146/annurev-soc-071913-043455

Firnkorn, J., Müller, M.: Selling mobility instead of cars: new business strategies of automakers and the impact on private vehicle holding. Business Strategy and the Environment **21**(4), 264-280 (2012). doi:DOI: 10.1002/bse.738

Greene, W.H., Hensher, D.A.: Modeling Ordered Choices - A Primer. Cambridge University Press, Cambridge, UK (2010)

Handy, S., Cao, X., Mokhtarian, P.: Correlation or causality between the built environment and travel behavior? Evidence from Northern California. Transportation Research Part D: Transport and Environment **10**(6), 427-444 (2005). doi:DOI:10.1016/j.trd.2005.05.002

Heckman, J.J., Ichimura, H., Smith, J., Todd, P.: Sources of Selection Bias in Evaluating Social Programs: An Interpretation of Conventional Measures and Evidence on the Effectiveness of Matching as a Program Evaluation Method. Proceedings of the National Academy of Sciences **93**(23), 13416-13420 (1996)

Ho, D.E., Imai, K., King, G., Stuart, E.A.: Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference. Political Analysis **15**(3), 199-236 (2007)

Ho, D.E., Imai, K., King, G., Stuart, E.A.: MatchIt: Nonparametric Preprocessing for Parametric Causal Inference - R Package Version 2.4-21. In. (2013)

Imbens, G.W., Rubin, D.B.: Causal Inference in Statistics, Social, and Biomedical Sciences. Cambridge University Press, (2015)

Kline, R.: Assumptions in Structural Equation Modeling In: Hoyle, R.H. (ed.) Handbook of Structural Equation Modeling. Guilford Publications, New York, NY (2012)

Kline, R.B.: Principles and Practice of Structural Equation Modeling, Third ed. Guilford press, New York, NY (2011)

Kopp, J., Gerike, R., Axhausen, K.W.: Do sharing people behave differently? An empirical evaluation of the distinctive mobility patterns of free-floating car-sharing members. Transportation **42**(3), 449-469 (2015). doi:DOI: 10.1007/s11116-015-9606-1

Loose, W.: The State of European Car-Sharing. In., p. 129. Bundesverband CarSharing e. V., (2010)

Maddala, G.: Limited-dependent and qualitative variables in econometrics. vol. 3. Cambridge University Press, (1986)

Martin, E., Shaheen, S.A., Lidicker, J.: Impact of Carsharing on Household Vehicle Holdings. Transportation Research Record: Journal of the Transportation Research Board **2143**(1), 150-158 (2010). doi:DOI: http://dx.doi.org/10.3141/2143-19

McKelvey, R.D., Zavoina, W.: A statistical model for the analysis of ordinal level dependent variables. Journal of mathematical sociology **4**(1), 103-120 (1975)

Mishra, G.S., Clewlow, R.R., Mokhtarian, P.L., Widaman, K.F.: The effect of carsharing on vehicle holdings and travel behavior: A propensity score and causal mediation analysis of the San Francisco Bay Area. Research in Transportation Economics **52**, 46-55 (2015). doi:DOI: 10.1016/j.retrec.2015.10.010

Mokhtarian, P.L., Cao, X.: Examining the Impacts of Residential Self-selection on Travel Behavior: A Focus on Methodologies. Transportation Research Part B: Methodological **42**(3), 204-228 (2008). doi:DOI: 10.1016/j.trb.2007.07.006

Morgan, S.L., Winship, C.: Counterfactuals and Causal Inference, Second ed. Cambridge University Press, (2014)

Moulton, B.R.: An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units. The Review of Economics and Statistics, 334-338 (1990)

Muthén, B.: Applications of Causally Defined Direct and Indirect Effects in Mediation Analysis using SEM in Mplus. (2011)

Muthén, L.K., Muthén, B.O.: Mplus User's Guide Seventh Edition. In. Muthén & Muthén, Los Angeles, CA, (1998-2014)

Rubin, D.B.: Using propensity scores to help design observational studies: application to the tobacco litigation. Health Services and Outcomes Research Methodology **2**(3-4), 169-188 (2001). doi:DOI:10.1023/A:1020363010465

Rubin, D.B.: Comment: The Design and Analysis of Gold Standard Randomized Experiments. Journal of the American Statistical Association **103**(484), 1350-1356 (2008)

Shaheen, S.A., Cohen, A.P., Chung, M.S.: North American Carsharing - 10-Year Retrospective. Transportation Research Record: Journal of the Transportation Research Board **2110**(1), 35-44 (2009)

Sioui, L., Morency, C., Trépanier, M.: How Carsharing Affects the Travel Behavior of Households: A Case Study of Montréal, Canada. International Journal of Sustainable Transportation **7**(1), 52-69 (2013). doi:DOI:10.1080/15568318.2012.660109

Waddell, P., Bhat, C., Eluru, N., Wang, L., Pendyala, R.: Modeling interdependence in household residence and workplace choices. Transportation Research Record: Journal of the Transportation Research Board **2003-11** (2007). doi:http://dx.doi.org/10.3141/2003-11

Xie, Y.: Structural Equation Models for Ordinal Variables An Analysis of Occupational Destination. Sociological Methods & Research **17**(4), 325-352 (1989)

Xie, Y.: Population heterogeneity and causal inference. Proceedings of the National Academy of Sciences **110**(16), 6262-6268 (2013). doi:DOI: 10.1073/pnas.1303102110

PART-IV

Paper 3: Does telecommuting reduce travel, or not? Accounting for self-selection using the UK National Travel Survey

Manuscript in preparation

DOES TELECOMMUTING REDUCE TRAVEL, OR NOT? ACCOUNTING FOR SELF-SELECTION USING THE UK NATIONAL TRAVEL SURVEY

Gouri Shankar Mishra (corresponding author)
Institute of Transportation Studies, University of California, Davis
gsmishra@ucdavis.edu

Patricia L. Mokhtarian
School of Civil and Environmental Engineering, Georgia Institute of Technology
patmokh@gatech.edu

Abstract

We estimate the effect of the adoption of telecommuting on the travel behavior (specifically, annual commute and non-work related travel) of full-time employed respondents with a fixed work location outside home, using data from the annual UK National Travel Surveys for the years 2009 to 2013. To control for self-selection bias, we adopt two distinct procedures that both use instrumental variables – the Two-Stage Predictor Substitution (2SPS) and Control Function (CF) methods.

On average, telecommuters are observed to travel more than non-telecommuters. However, after accounting for the *observed* differences in *traits* and *tastes* between the two groups using a linear regression model, the differences fade to (nearly-) insignificant levels. Further control of self-selection bias arising from *unmeasured* differences in "relevant" characteristics leads to the conclusion that telecommuting has a substitution effect on both commute and non-work travel. The CF method identifies the nature of the self-selection bias – the net effect of unmeasured factors both encourages the adoption of telecommuting and leads to higher volumes of travel. This correlation of unobserved variables is stronger for telecommuters than non-telecommuters. The effect on non-work travel is imprecisely estimated, i.e. having large standard errors. This suggests that the telecommuting effect may be highly variable, with both substitution and complementarity effects in play albeit in different proportions for different people.

Keywords: *Program evaluation, omitted variable bias, treatment effects, selection on unobservables, endogeneity, Control Functions.*

1. Introduction

The effects of telecommuting on travel behavior have been studied and debated extensively in the literature. For the most part, these studies have relied on disaggregate data related to specific telecommuting programs, obtained through the administration of travel diaries (generally to telecommuters as well as a control group of non-telecommuters, and sometimes to household members of both groups) before and after the adoption of telecommuting – i.e. using panel data (Mokhtarian, Handy and Salomon 1995; Andreev, Salomon and Pliskin 2010). The samples involved were relatively small, and the purpose of the travel diaries was explicitly to evaluate the impacts of telecommuting. By implicitly assuming that the unobserved influencers of both telecommuting adoption as well as travel behavior (such as attitudes or lifestyles) do not change over the course of the evaluation (an assumption which may or may not be true), the studies estimate the effect of telecommuting by comparing travel behavior before and after telecommuting adoption relative to that of a control group.

As telecommuting became more widespread, regional and national travel surveys began to measure it. These surveys also involved travel diaries, but differed from the earlier studies in three key respects. First, the samples were much larger in size and geographic scope. Second, the study purpose stated to respondents was quite general (on the order of "to better understand travel patterns"), rather than specific to a telecommuting application. Third, the samples were cross-sectional rather than longitudinal. The first two distinctions constitute advantages relative to the panel studies (in the second case because it would have reduced or eliminated any response and non-response biases associated with telecommuting). A number of studies have used these cross-sectional data to study the influencers of telecommuting adoption (Popuri and Bhat 2003; Singh et al. 2013; Walls, Safirova and Jiang 2007). However, only a few studies, most of them published in the recent past, have capitalized on these datasets to estimate effects of telecommuting on travel behavior (He and Hu 2015; Kim, Choo and Mokhtarian 2015; Zhu 2013; Zhu and Mason 2014).

Interestingly, studies relying on cross-sectional data have tended to find opposite impacts of telecommuting on travel compared to the studies relying on panel data. With respect to the panel studies, Kim, Choo, Mokhtarian (2015) and Andreev, Salomon, Pliskin (2010) conclude that most of them find a substitution effect i.e. telecommuting leads to less travel, and the results are usually substantial and statistically significant. These conclusions hold for both commute as well as non-commute travel. The cross-sectional studies, on the other hand, have found a complementary effect, with telecommuting leading to more, rather than less, travel (Zhu 2013; Zhu and Mason 2014) or at least associated with more travel even after controlling for various observed differences between telecommuters and non-telecommuters (He and Hu 2015; Kim, Choo and Mokhtarian 2015). It is possible that this discrepancy arises due to the difficulty in controlling for what is known in the econometric literature as "self-selection bias" or "omitted variable bias" in cross-sectional observational studies (Kim, Choo and Mokhtarian 2015).

Such bias arises because not all variables that influence respondents' travel behavior *and* telecommuting adoption decision are measured. For example, cross-sectional travel surveys like the US National Household Travel Survey (NHTS) or the UK National Travel Survey (UK-NTS) collect a wide range of demographic characteristics but little or no information on attitudes towards such issues as the environment, driving, transit use, and urban living. However, such attitudes have been shown to considerably improve the ability to explain people's decisions regarding travel (Handy, Cao and Mokhtarian 2005), as well as the adoption of telecommuting (Mokhtarian and Salomon 1997). In the causal inference literature, the bias is attributed to the mechanism by which respondents adopt the treatment or program (in our case, telecommuting). In randomized experiments, the gold standard for estimating program effects, respondents are assigned to the program or no-program (control) groups on a random basis in order to minimize any systematic observed or unobserved differences. Contrary to that standard, assignment in an observational sample is a non-random choice-based process. Respondents are likely optimizing their overall utility from travel as well as from activities that are influenced by travel decisions while deciding on the adoption of telecommuting, and thus those who choose to telecommute are likely to differ systematically from those who choose not to do so. When the factors influencing this nonrandom decision-making process are *measured*, to a large extent they can be controlled for by including them as observed explanatory variables in the model. When they are *unmeasured*, however, the causal effects of telecommuting are conflated with self-selection bias in the observed association between telecommuting and travel behavior.

In this paper, we present two alternative approaches to control self-selection bias and identify the effect of telecommuting using the UK-NTS, a large cross-sectional dataset. Both approaches depend on the availability of what are known as instrumental variables (IV) – exogenous variables necessary to partition the variation in *d* into two parts, separating the part that is uncorrelated with the error ("good" variation) from the part correlated with the error ("bad" variation) and leading to endogeneity. We measure the nature of self-selection bias, and test for differences in direction and magnitude of this bias between telecommuters and non-telecommuters.

2. Sources Of Endogeneity & Self-Selection Bias

The effect of telecommuting (d) on some travel behavior outcome y for an individual *i* may be expressed using the following equation:

(1A) $$y_i = \alpha + \boldsymbol{z_{i,M}}\boldsymbol{\beta_M} + d_i\tau + \boldsymbol{z_{i,U}}\boldsymbol{\beta_U}\ .$$

In (1A), $\boldsymbol{z_{i,M}} = [z_{i,M1} \dots z_{i,MK}]$ is a 1xK vector of exogenous variables representing the observable characteristics of individual *i* (variously referred to as observed confounders or covariates, see Table 1). Correspondingly, $\boldsymbol{\beta_M}$ is a Kx1 vector of parameters. d_i reflects the telecommuting status of respondent *i,* and takes on a value of 1 if the respondent telecommutes and 0 otherwise (we define telecommuting more precisely at the end of this section). τ represents the effect of telecommuting and is the parameter of interest. In the causal inference literature, *d* may be referred to as the treatment indicator. $\boldsymbol{z_{i,U}}$ is a J-dimensional vector (J is unknown) of unobserved variables which are expected to influence the travel behavior outcome y_i and $\boldsymbol{\beta_U}$ is the conformable vector of coefficients of those variables. The subscripts "M" and "U" stand for measured (and thus observed) and unmeasured (and thus unobserved) respectively.

Given that the elements of $\boldsymbol{z_{i,U}}$ are unmeasured, (1A) cannot be estimated in its entirety. Instead, we may estimate the equation

(1B) $y_i = \alpha + \boldsymbol{z_{i,M}\beta_M} + d_i\tau + \varepsilon_i$,

where ε_i constitutes $\boldsymbol{z_{i,U}\beta_U}$, the net influence of all unobserved variables on y_i. Potential for endogeneity arises when $\boldsymbol{z_{i,U}}$ also influences d_i, i.e. when d_i may be expressed as

(2A) $$d_i = \begin{cases} 1, if\ d_i^* = \boldsymbol{z_i\pi} + \boldsymbol{z_{i,U}\pi_U} > 0 \\ 0,\ \text{otherwise} \end{cases}.$$

In (2A), which we refer to as the selection equation, d_i^* may be construed as the propensity to telecommute or the utility derived from telecommuting. If the latent utility is above a threshold (0 for convenience), then a person adopts telecommuting. Here, $\boldsymbol{z_i}$ is a vector of observed exogenous variables and consists of $\boldsymbol{z_{i,M}}$ and a set of IVs $\boldsymbol{z_{i,IV}}$ (we discuss this further in Section 4). Again, given that $\boldsymbol{z_{i,U}}$ is unmeasured, we have to rewrite (2A) as the following (indicating that $\boldsymbol{z_{i,U}\pi_U}$ comprises the error term v_i):

(2B) $$d_i = \begin{cases} 1, if\ d_i^* = \boldsymbol{z_i\pi} + v_i > 0 \\ 0,\ \text{otherwise} \end{cases}.$$

The above setup highlights the endogeneity problem arising from the exclusion of $\boldsymbol{z_{i,U}}$ in (1A) and (1B), which leads to $E(d\varepsilon) \neq 0$ in violation of conventional assumptions. Given that the errors ε_i and v_i both contain $\boldsymbol{z_{i,U}}$, i.e. $\varepsilon_i = \boldsymbol{z_{i,U}\beta_U}$ and $v_i = \boldsymbol{z_{i,U}\pi_U}$, an alternate way to describe the endogeneity problem is that the errors are correlated: $corr(\varepsilon, v) \equiv \rho \neq 0$. Under these circumstances, d_i is referred to as an endogenous explanatory variable (EEV) in the econometric literature. This endogeneity from common omitted variables is referred to as "self-selection bias" or "selection bias due to unobservables" in the causal inference literature. The result of d_i being an EEV is that the estimate of its coefficient (τ, the influence of telecommuting) based on the standard ordinary least squares (OLS) approach will be biased (i.e. $E[\hat{\tau}_{OLS}] \neq \tau$).

In the above discussion, we assumed that the effects of $\boldsymbol{z_{i,M}}$ and $\boldsymbol{z_{i,U}}$ on *y* are the same regardless of whether person *i* is a telecommuter or a non-telecommuter. We can relax this assumption and allow for differences between the two states. To motivate this distinction, we re-write equation (1A) to have separate outcome equations for the two regimes, i.e. telecommuting adopted (d_i=1) and telecommuting not adopted (d_i=0), and allow the coefficients $\boldsymbol{\beta_M}$ and $\boldsymbol{\beta_U}$ to differ between the two regimes:

$$\text{(1C)} \quad \begin{cases} y_{i,d_i=1} = \alpha_1 + \boldsymbol{z_{i,M}\beta_{M1}} + \boldsymbol{z_{i,U}\beta_{U1}} \text{ when } d_i^* > 0 \quad \text{(telecommuting adopted)} \\ y_{i,d_i=0} = \alpha_0 + \boldsymbol{z_{i,M}\beta_{M0}} + \boldsymbol{z_{i,U}\beta_{U0}} \text{ when } d_i^* \leq 0 \text{ (telecommuting not adopted)} \end{cases}$$

$$\text{(1D)} \quad \begin{cases} y_{i,d_i=1} = \alpha_1 + \boldsymbol{z_{i,M}\beta_{M1}} + \varepsilon_{i,1} \text{ when } d_i^* > 0 \quad \text{(telecommuting adopted)} \\ y_{i,d_i=0} = \alpha_0 + \boldsymbol{z_{i,M}\beta_{M0}} + \varepsilon_{i,0} \text{ when } d_i^* \leq 0 \text{ (telecommuting not adopted)} \end{cases}$$

In this most general form, the model (1C) or (1D) is alternatively referred to in the literature as endogenous switching regression, a "mover/stayer" or Roy model, or a two-outcome version of the standard "Heckit" model (Mokhtarian and van Herick 2016); or even a "Tobit-5" model (Toomet and Henningsen 2008). (1A) may be considered a special case of (1C) where $\boldsymbol{\beta_{M1}} = \boldsymbol{\beta_{M0}}$, $\boldsymbol{\beta_{U1}} = \boldsymbol{\beta_{U0}}$, $\alpha = \alpha_0$, and $\tau = \alpha_1 - \alpha_0$. An assumption of inequality in $\boldsymbol{\beta_M}$ for the two groups implies *heterogeneity in observed tastes* (Section 4.4); inequality in $\boldsymbol{\beta_U}$ implies *heterogeneity in unobserved tastes* (Section 4.5). As an example of the latter, suppose that unobserved factors (such as a dislike of travel) that increase the propensity to telecommute also decrease the overall levels of travel (the outcome of interest in this study) for both groups. If the impact on travel of a given degree of dislike of travel is greater for telecommuters than for non-telecommuters, the strength of the correlation between unobserved variables influencing telecommuting adoption and travel will tend to be higher for the former group than for the latter.

It is important to understand that, *in principle*, the d=1 and d=0 equations of (1C) and (1D) are not intended to apply to *two different groups of people*, but to the *same people in two different states*: telecommuting and not telecommuting. Otherwise, it would not make sense to talk about an effect of treatment (telecommuting), which necessarily implies that a given person is untreated at one point, and then becomes treated. Thus, in principle, we need two sets of labels: one to denote *group membership* (those who have been designated to be treated versus those who have been designated not to be treated), and the other to denote possible *states* or *conditions*: treated (telecommuting) versus untreated (not telecommuting). Members of either group could hypothetically be observed in either condition. But because with cross-sectional data we only observe treatment group members (telecommuters) in the treatment condition (telecommuting) and others (non-telecommuters) in the untreated (non-telecommuting) condition, it is (unfortunately) common to confound the two needed sets of labels.

Furthermore, because we do not observe cases in both conditions, we necessarily use the telecommuters to estimate the first outcome equation of each pair, and the non-telecommuters to estimate the second outcome equation, in both cases correcting for the bias arising from the correlation of d with ε. For this reason, the term "heterogeneity" is often used as if referring to two different groups of people ("telecommuters' coefficients differ from those of non-telecommuters'") rather than the same people in two different conditions ("the explanatory variables have a different influence on individual i's travel behavior during a period when she has adopted telecommuting than during a period when she has not"). But it is the latter situation that is reflected when the model is used to estimate the average treatment effect: for each case, the estimated outcome $\hat{y}_i$ is computed under the non-telecommuting condition ($\hat{y}_{i,d_i=0} = \hat{\alpha}_0 + \boldsymbol{z_{i,M}}\widehat{\boldsymbol{\beta}}_{\boldsymbol{M0}}$), and subtracted from the estimated outcome under the telecommuting condition

($\hat{y}_{i,d_i=1} = \hat{\alpha}_1 + \boldsymbol{z_{i,M}}\widehat{\boldsymbol{\beta}}_{\boldsymbol{M1}}$). The analysis approach has corrected the $\widehat{\boldsymbol{\beta}}_{\boldsymbol{M}}$s so that they are consistent estimators of the true $\boldsymbol{\beta_M}$s under each condition.

In the most general case (as reflected in equations (1D)), we not only allow for $\boldsymbol{\beta_{M1}} \neq \boldsymbol{\beta_{M0}}$, but also for $\boldsymbol{\beta_{U1}} \neq \boldsymbol{\beta_{U0}}$. This has two consequences: (a) the variances of ε_1 and ε_0 are different (i.e., we assume $\varepsilon_1 \sim N(0, \sigma_1^2)$ and $\varepsilon_0 \sim N(0, \sigma_0^2)$); and (b) their correlations with υ are different, potentially in sign as well as magnitude (i.e. we assume $\rho_{d=1} \equiv corr(\varepsilon_1, \upsilon) \neq corr(\varepsilon_0, \upsilon) \equiv \rho_{d=0}$). Some applications allow for $\boldsymbol{\beta_{M1}} \neq \boldsymbol{\beta_{M0}}$ but constrain $\sigma_1^2 = \sigma_0^2$ and $\rho_{d=1} = \rho_{d=0}$, so in Section 4.4 we impose those constraints, and relax them in Section 4.5.

The assumption $\boldsymbol{\beta_{U1}} \neq \boldsymbol{\beta_{U0}}$ constitutes one form of *heterogeneity in self-selection*, namely (as mentioned above) *heterogeneity in unobserved tastes.* However, given that each outcome equation is estimated using a different group of people (current telecommuters and current non-telecommuters, respectively), heterogeneity in self-selection bias can occur in a second way as well: *heterogeneity in distribution of unobserved traits*. This arises when there are systematic differences in unobserved characteristics between telecommuters and non-telecommuters, i.e. when the distributions of $\boldsymbol{z_{i,U}}$ differ between the two groups. This will also result in $\sigma_1^2 \neq \sigma_0^2$ and $\rho_{d=1} \neq \rho_{d=0}$. As an example, suppose that compared to non-telecommuters, telecommuters (1) are, on average, more risk-taking, variety-seeking, and adventurous (i.e. have different means), and (2) have less diversity on these traits (different variances). These latent traits could potentially lead to greater non-work travel *even before adopting telecommuting*. In addition, to the extent that they moderate the influence of any unmeasured covariates, this heterogeneity in distribution of unobserved traits can also contribute to heterogeneity in unobserved tastes. It is possible that the two groups differ not only in terms of magnitude of ρ, but also in terms of direction. For example, telecommuters may embrace a high-activity lifestyle while non-telecommuters may tend to be more social, more inclined to like travel for its own sake, have a less work-oriented lifestyle – factors which are *also* likely to be associated with greater non-work travel.

In this paper, we identify the nature of self-selection and how it differs between the telecommuters and non-telecommuters (relative magnitudes and signs of $\rho_{d=0}$ & $\rho_{d=1}$).

Endogeneity and correlation between d and ε may also arise from two other sources of relevance to this study. First, there may be reverse causality or simultaneity bias. Reverse causality arises "…when an "explanatory" variable is simultaneously a function of the "dependent" variable it is supposed to explain – that is, when one variable is both a cause and an effect of another…" (Mokhtarian and Cao 2008, p. 206). In this study, the effect of interest is the influence of telecommuting on the level of travel. However, high levels of travel even before the adoption of telecommuting may influence the adoption of telecommuting, leading to endogeneity. Second, we may incorrectly induce bias by conditioning on variables that are either outcomes of *d* or *y*. This is called over-conditioning or endogenous selection bias (Elwert and Winship 2014; Angrist and Pischke 2008). We describe both biases in Section 4 and discuss our approach to control (at least partially) these biases.

3. Data

In this study, we estimate the effects of telecommuting using the UK National Travel Survey, a survey of approximately 15,000 individuals from 7,000 households in Great Britain – England, Scotland, and Wales – conducted annually since 2002. Since 2013, the survey's geographic coverage was limited to England only. The cross-sectional survey involves a face-to-face interview followed by a seven-day travel diary for all members of the household. Starting in 2009, all employed (but not home-based) respondents indicated the frequency with which they work from home instead of commuting to a workplace – from less than once a year or never to three or more times a week. This excludes days when the respondent worked only partly from home (including evenings and weekends). Respondents who worked from home instead of commuting to a workplace at least once a month are categorized as telecommuters in this study ($d = 1$).

Given large changes in the survey questionnaire in 2009, especially in those sections involving telecommuting, we excluded data collected before 2009. Further, a few critical variables were not collected in the 2014 survey; hence we only considered data collected through 2013.

We restrict our sample to full-time employed respondents who have a fixed work location outside home. Based on the above filtering criteria, we have 11,471 non-telecommuters and 1,475 telecommuters. The travel behavior outcomes considered in this study are (i) total annualized travel not including commute or business-related travel (non-work), and (ii) total annualized commute travel. Both person-miles traveled (PMT) and vehicle-miles traveled (VMT) are considered. To calculate the annualized distances, we summed up the travel distances in the seven-day travel diary disaggregated by trip purposes and travel mode; and then annualized them using methods discussed in UK Department for Transport (2015b). Per UK Department for Transport (2015a), around 16% of commute trips are chained – hence total commute travel may be under-measured because of being allocated to non-commute trip purposes. More importantly, it is possible that the adoption of trip chaining may be different between telecommuters and non-telecommuters; we make a simple assumption that the adoption of trip chaining does not differ between the two groups. We will relax this assumption in future analysis.

Figure 1 shows the distribution of annual travel disaggregated by mode (all modes and car only) and trip purpose (non-work, commute and business-only). The figure indicates that, on average, telecommuters travel more than non-telecommuters across all three aggregated trip purposes. For both groups, the median business-only distance traveled is zero, but 81% of the non-telecommuters did not report any business trips in their one-week travel diary, while the corresponding number for telecommuters is only 55%.

Figure 1. Observed annual travel disaggregated by trip purpose and travel mode

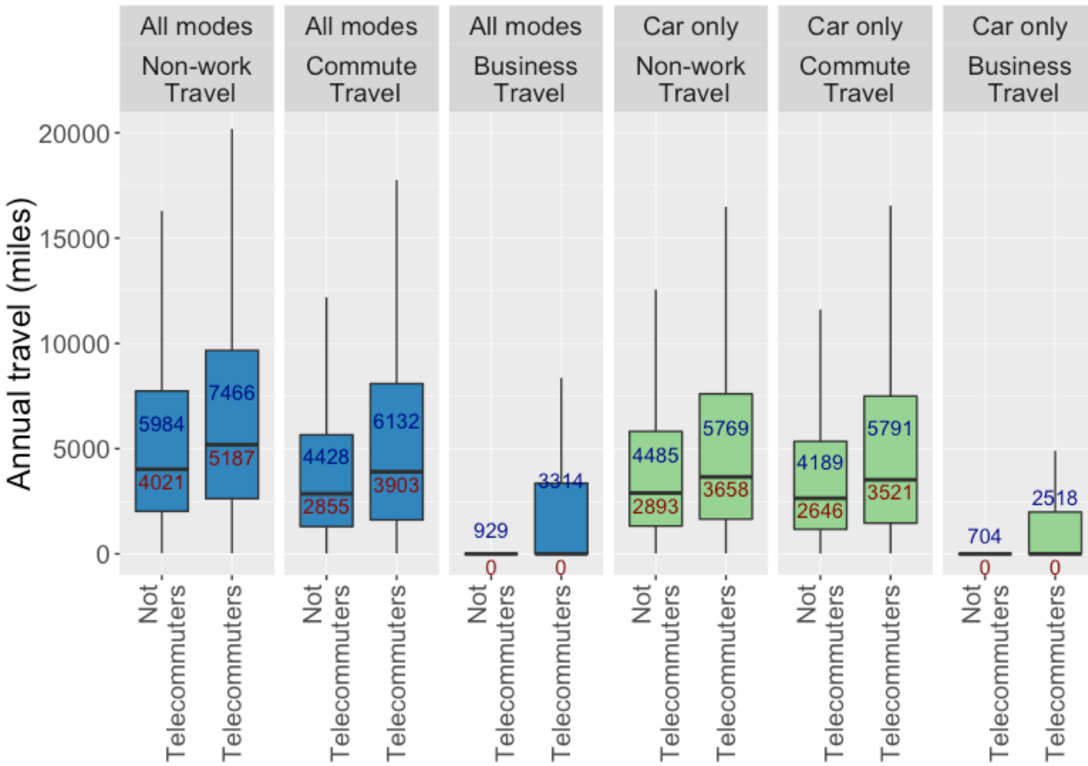

Note: The top number in blue font gives the mean, the bottom number in dark red gives the median. Sample size: 11,471 non-telecommuters and 1,475 telecommuters.

The graph confirms conclusions that telecommuters travel more than non-telecommuters. However, the higher levels of travel may not be attributed to telecommuting, at least not entirely. The difference may be explained partly by differences in observed and unobserved characteristics and tastes between the two groups. This study aims to disentangle the effect of these factors, and identify the effect of telecommuting, using econometric procedures.

Table 1. Key variables used in this study

Variable name	Description
Travel behavior outcome and telecommuting (treatment) indicator	
Annual all travel except business travel (miles, $y1$ & $y1_{car}$) Annual commute travel (miles, $y2$ & $y2_{car}$) Non-work annual travel (miles, $y3$ & $y3_{car}$)	$y1$ captures all travel undertaken by an individual except business trips. $y2$ includes only commute travel. $y3$ aggregates all travel with trip purposes other than commute and business, hence referred to as non-work related). The variables summarize travel across all modes. ys with subscript "car" represent travel in car/van only. All variables are natural log transformed. Estimates are based on a seven-day travel diary. All trips, including short walk trips, were recorded, and the respondent provided details of origin and destination, purpose, mode, distance travelled, time, vehicles used, etc. The weekly record was annualized based on methods suggested by UK Department for Transport (2015b).
Telecommuter flag (d)	1 if a person telecommutes once a month or more frequently, 0 otherwise. Indicator of treatment or program participation.
Observed confounders – Individual and HH characteristics	
Frequency of business trips	Number of annual business trips. As in the case of travel distance, the weekly record of business trips is annualized.
Basic respondent demographics	Multiple variables: Age, ethnicity, gender, marital status (married or cohabitating with partner), education (does respondent have a college degree), employment status (full time or not), and student status (is respondent enrolled in school).
Proxy response	Were the interview questions asked face to face (1) or by proxy, i.e. another HH member answered for the respondent (0)
Commute mode	Primary mode of commuting to work. Five aggregated modes including car/van (as either driver or passenger), transit (includes rail, bus, and tram), bike, walk & others.
Respondent's socioeconomic group	The National Statistics Socioeconomic Classification (NS-SEC) adopted by the UK Office of National Statistics (Rose, Pevalin and O'Reilly 2005). Employed respondents fall under four aggregate categories: (1) Managerial & professional occupations, (2) Intermediate occupations and small employers, (3) Routine and manual occupations, (4) Not classified, including students. Further, we disaggregate category (1) into three sub-categories: (1a) Large employers and higher managerial occupations, (1b) Higher professional occupations, (1c) Lower managerial and professional occupations.
Industry classification	Standard Industrial Classification (2007 categorization). The 25+ categories in the original dataset were aggregated to 14 categories for our analysis. The industry classification for the 2009 survey data used the older 1992 categorization, and had to be reclassified.
Air travel frequency	An ordinal variable giving frequency of internal air flights within Great Britain. From 3 or more flights a week (1) to less than once a year (7)
Household income	HH income semi-deciles. For any survey year, the HHs were classified into semi-deciles independent of responses from other survey years. We aggregated the data into a single income semi-deciles variable.
HH structure	Two variables treated as continuous: (1) Number of adults in the HH, (2) Number of children in the HH. Also, two binary variables: (1) Is there a

	telecommuter in HH other than respondent, (ii) Is there an employed non-telecommuter in the HH other than respondent.
Vehicle & bicycle holdings	Two count variables indicating number of cars & vans, and bicycles owned or leased by the household.
Built environment – central workplace	A binary variable equal to 1 if respondent's workplace is located in a metropolitan central area, 0 otherwise
Built environment – residential area population density	A continuous variable giving population density of residential neighborhood. The density is calculated at the level of Primary Sampling Unit (PSU) identified based on zip codes (postal codes). The sample was drawn each survey year by first selecting the PSUs (around 680) and then selecting addresses within each PSU. Half the PSUs were drawn from the previous year, and another half were randomly selected. Hence, in our sample of data from 2009-2013, we have representation from around 2000 PSUs (Taylor et al. 2013)
Other Covariates	
Survey year	Dummy codes for each year of survey to capture the declining trend in per capita travel in the UK
Economic prospects	Three variables summarizing the annualized GDP growth rate (compared to same quarter in previous year) for the quarter of respondent's response and prior two quarters. A fourth variable captures the variability in daily returns of the stock market index – FTSE-1000 – in the same quarter. We included these variables to test the hypothesis that telecommuting adoption decisions are influenced by economic prospects and perceptions of job security.
Potential Instrumental Variables	
Internet access[1]	1 if respondent's HH has access to "broadband" internet, 0 otherwise (includes no access as well as dial-up connection). Based on arguments in Section 4, this variable was not adopted as an instrument but was included as a covariate.
Potential to telecommute	How much of the respondent's work can be "theoretically" done at home? Assuming availability of necessary equipment at home and not considering employer's policies.

[1]Based on arguments in Section 4.7, this variable does not meet the criteria for a valid instrument.

4. Methods

In this section, we discuss the two procedures to address self-selection bias and identify causal effects. Readers may refer to the following publications for more detailed discussions of the methods (Angrist and Pischke 2008; Cerulli 2015; StataCorp 2015; Wooldridge 2010, 2015; Madalla 1983).

4.1 Overview

The first step of both methods included in this study involves estimation of the propensity to telecommute as a function of vector $\boldsymbol{z}$ (Eq. 2B). $\boldsymbol{z}$ includes all observed exogeneous variables $\boldsymbol{z_M}$ included in the outcome model (1B) and (1D), and a set of L (instrumental) variables that influence *d* but not *y* i.e. $\boldsymbol{z} = [\boldsymbol{z_M}, \boldsymbol{z_{IV}}]$. As a result d^* in (2B) may be represented as $d^* = \boldsymbol{z_M}\boldsymbol{\pi_M} + \boldsymbol{z_{IV}}\boldsymbol{\pi_{IV}} + \upsilon$.

The IVs $\boldsymbol{z_{IV}} = [z_{IV1} \dots z_{IVL}]$ should meet two specific conditions: (A1) the IV(s) explain some of the variation in *d,* the decisions to adopt telecommuting or not (instrument relevance), and (A2) the sole channel through which the IV(s) affect *y* is through *d* (exclusion restriction or instrument exogeneity). If met, the second condition ensures that the IVs are uncorrelated with the unobserved variables $\boldsymbol{z_U}$.

Assumption A1 for relevant IVs may be restated as $\boldsymbol{\pi_{IV}} \neq \boldsymbol{0}.$ This assumption confirms that the IVs explain some of the variation in *d.* To be suitable and provide consistent estimates of τ, the IVs should explain a "sufficiently" large portion of the variation in *d.* Both aspects of assumption A1 can be tested. In the estimation of equation (2C), are the coefficients $\boldsymbol{\pi_{IV}}$ statistically significant? Further, is the difference in pseudo R-squared (or any other measure of goodness-of-fit) of (2C) estimated with and without $\boldsymbol{z_{IV}}$, large enough?

Unlike A1, assumption A2 cannot be formally tested and must be qualitatively argued based on substantive knowledge and literature review. In addition to the assumptions about IVs, the alternate methods adopted in this paper make two additional assumptions: (A3) Vector $\boldsymbol{z} = [\boldsymbol{z_M}, \boldsymbol{z_{IV}}]$ is exogenous, which may be expressed in terms of the following $2\times(K+L)$ orthogonality or zero-covariance conditions: $E(z_j\varepsilon) = 0$, and $E(z_j\nu) = 0,\ \ j = 1,2, \dots (K+L)$. (A4) υ in (2B) is assumed to be normally distributed, allowing us to estimate the propensities to telecommute using a probit regression:

$$\text{(2C)} \quad P(d = 1|\boldsymbol{z_M}, \boldsymbol{z_{IV}}) = \Phi[\boldsymbol{z_M}\boldsymbol{\pi_M} + \boldsymbol{z_{IV}}\boldsymbol{\pi_{IV}}] = d_{P1}.$$

As is standard in probit regressions, the variance of υ is fixed at 1 for identification.

In both methods adopted in the study, the propensity to telecommute is first estimated (first stage reduced form, 2C), followed by estimation of the outcome equation(s). In *Two-Stage Predictor Substitution* (2SPS) the endogenous variable *d* is substituted with a "linearized version" of estimated propensities $\hat{d}_{P1}$. The estimated propensities contain only the "good variation" in *d* accounted for by the IVs (and the observed variables in common with *y*, $\boldsymbol{z_M}$) and thus remove its correlation with error term ε. In the second method – the Control Functions (CF) method – endogeneity is addressed by including "generalized residuals" for each respondent estimated using the probabilities from the selection equation. These generalized residuals contain information about the "bad" variation, which when included in the outcome equation (1B) make the EEV appropriately exogenous.

The 2SPS and CF methods, in their simplest and least flexible form (Sections 4.2 and 4.3), assume homogeneity in both observed tastes as well as nature and magnitude of selection bias in both the regimes (as encapsulated in Eq. 1B). These restrictive versions are presented for pedagogical reasons only; relaxing the assumption of homogeneity in observed tastes (i.e. allowing $\boldsymbol{\beta_{M1}} \neq \boldsymbol{\beta_{M0}}$) as done in Sections 4.4 & 4.5 should be the adopted approach in any study. In Section 4.4, we continue to assume homogeneity in self-selection bias i.e. $\rho_{d=1} = \rho_{d=0} = \rho$. In all of these scenarios, the variance-covariance matrix of the error terms ε and υ may be written as $\begin{bmatrix} \sigma^2 & \rho\sigma \\ \rho\sigma & 1 \end{bmatrix}$ where σ^2 is the variance of ε and the variance of υ is fixed at 1.

In Section 4.5, we adopt a fully flexible CF model that relaxes the assumption of homogeneity in nature and magnitude of self-selection bias between the two regimes, i.e. where we allow $\rho_{d=1} \neq \rho_{d=0}$. We should note that the flexible form is restricted to the CF method only; to our knowledge, there is no equivalent 2SPS method. The outcome equation under this flexible form is captured by Eq. (1D).

We should reiterate that $\boldsymbol{z_M}$ is a strict subset of $\boldsymbol{z}$**.** This implies that all components of $\boldsymbol{z_M}$ should be included in estimating the propensity to telecommute. Omission of one of more components of $\boldsymbol{z_M}$ will lead to an inconsistent estimate of the τ. The requirement for $\boldsymbol{z_M}$ to be a strict subset of $\boldsymbol{z}$ also implies that $\boldsymbol{z_{IV}}$ is not a null vector and its dimension L is at least equal to 1. This is necessary because the "exogenous variation" induced by excluded instrumental variables provides separate variation in the estimated propensities (necessary for the 2SPS method) and the generalized residuals (necessary for the CF method). Both methods rely on this exogenous variation to identify causal effects. We should note that this runs counter to the common practice in literature across disciplines where the Heckit model has been sometimes estimated without any IVs. In the absence of IVs, the outcome equation regression will suffer from collinearity leading to inflated standard errors and hence a greater potential for Type II errors (Wooldridge 2015; Bushway, Johnson and Slocum 2007).

4.2 2SPS with homogeneous taste & self-selection bias assumption

2SPS is an adaptation of the 2SLS instrumental variables method (reader should note that we make a distinction between the instrumental variables *per se* ($\boldsymbol{z_{IV}}$), and the "IV method" which is one particular approach using IVs to control endogeneity). The standard IV method is applicable only in linear models, where the endogenous variable is continuous and appears in linear form in the structural model. When the linearity assumptions are met, the estimated EEVs from the first stage OLS may be substituted in the structural model to control for endogeneity. In this paper, the linearity assumption is violated because the endogenous variable *d* is a dichotomous variable and (2B) is a non-linear model.

An incorrectly adopted alternate in the literature to address this non-linearity is to substitute *d* in equation 1B with the fitted values from the probit first stage ($\hat{d}_{P1}$). Although in the case of OLS regression the residual ($d - \hat{d}_{P1}$) will be independent of d, the independence does not hold here because (2C) is a non-linear probit model. As a result, the second-stage structural equation $y = \alpha + \boldsymbol{z_M}\boldsymbol{\beta_M} + \hat{d}_{P1}\tau + \left[\varepsilon + \tau(d - \hat{d}_{P1})\right]$ may not produce consistent estimates of the treatment effect and is referred to as a "forbidden regression" (Angrist and Pischke 2008, Section 4.6).

Instead, the 2SPS method treats the estimated propensities $\hat{d}_{P1}$ as an additional instrument for *d*. Using this instrument and $\boldsymbol{z_{IV}}$, an estimation process that mimics the 2SLS / IV method may be applied. An OLS regression of *d* on $\boldsymbol{z_M}$ and instrument $\hat{d}_{P1}$ leads to estimates of $\hat{d}_{P2}$, which when substituted for *d* in the structural equation gives:

(3) $$y = \alpha + \boldsymbol{z_M}\boldsymbol{\beta_M} + \hat{d}_{P2}\tau + \left[\varepsilon + \tau(d - \hat{d}_{P2})\right].$$

Since (3) is estimated via OLS, the variable $\hat{d}_{P2}$ is independent of $(d - \hat{d}_{P2})$ and therefore presumed independent of the error term $\left[\varepsilon + \tau(d - \hat{d}_{P2})\right]$, thereby addressing the endogeneity problem. Here, the estimate $\hat{\tau}_{2SPS}$ is more likely to be a consistent estimate of the parameter of interest τ. Thus, the estimation follows three stages:

(i) Estimate $\hat{d}_{i,P1}$ from equation (2C), using $\boldsymbol{z_{i,M}}$ and $\boldsymbol{z_{i,IV}}$ for respondent $i, i = 1, \dots, N$.

(ii) Run the following **OLS** regression and obtain $\hat{d}_{i,P2}$:

d_i on 1, $\hat{d}_{i,P1}$, $\boldsymbol{z_{i,M}}$, $\quad i = 1, \dots, N.$

(iii) Run the OLS regression:

y_i on 1, $\hat{d}_{i,P2}$ and $\boldsymbol{z_{i,M}}$, $\quad i = 1, \dots, N.$

In the Supporting Information (SI), we provide the "forbidden regression" results where $\hat{d}_{P1}$ substitutes for *d* in the outcome equation thereby (improperly) eliminating step (iii) above.

4.3 CF with homogeneous taste & self-selection bias assumptions

In the control function approach, a CF estimated from (2B) is included in the structural equation as an additional "explanatory variable".

(4) $$E(y|\boldsymbol{z_M}, d) = \alpha + \boldsymbol{z_M}\boldsymbol{\beta_M} + d\tau + \rho\sigma[d\lambda(\boldsymbol{z\pi}) - (1-d)\lambda(-\boldsymbol{z\pi})]$$

where $\lambda(\cdot) = {}^{\phi(\cdot)}\!/_{\Phi(\cdot)}$ is the well-known Inverse-Mills Ratio (IMR) which is the ratio of the standard normal probability density function ϕ to the standard normal cumulative distribution function Φ, where the arguments of both functions are the observed utility of the selection equation (or its negative). When a person telecommutes (d=1), the control function added to the structural equation is $\phi(\boldsymbol{z\pi})/\Phi(\boldsymbol{z\pi})$; else it is $-\phi(\boldsymbol{z\pi})/\Phi(-\boldsymbol{z\pi})$. The function $r(d, \boldsymbol{z\pi}) \equiv d\lambda(\boldsymbol{z\pi}) - (1-d)\lambda(-\boldsymbol{z\pi})$ is sometimes called the "generalized residual" term because it has a mean of zero conditional on **z**. Its coefficient is the product of error correlation (ρ) and error covariance ($\sigma \times 1 = \sigma$). We refer to the estimated coefficient of d as $\hat{\tau}_{CF}$.

The method can be estimated in a multi-step procedure similar to 2SLS:

(i) Estimate $\hat{d}_{i,P1}$ from equation (2B), using $\boldsymbol{z_{i,M}}$ and $\boldsymbol{z_{i,IV}}$ for respondent $i, i = 1, \dots, N$.

(ii) Estimate the generalized residuals,

$$\hat{r}_i \equiv d_i\lambda(\boldsymbol{z_i}\widehat{\boldsymbol{\pi}}) - (1-d_i)\lambda(-\boldsymbol{z_i}\widehat{\boldsymbol{\pi}}), \qquad i = 1, \dots, N.$$

(iii) Run the following OLS regression:

y_i on 1, d_i, $\boldsymbol{z_{i,M}}$, $\hat{r}_i$, $\qquad i = 1, \dots, N$.

4.4 Observable heterogeneity in tastes but homogeneous self-selection bias (Methods: H.2SPS and H.CF)

To relax our assumption of homogeneity in tastes, we consider the structural model in (1D) where the outcome model was defined separately for the two regimes. The equation is reproduced here for convenience:

(1D) $$\begin{cases} y_{d=1} = \alpha_1 + \boldsymbol{z_M}\boldsymbol{\beta_{M1}} + \varepsilon_1 \text{ when } d^* > 0 \quad \text{(telecommuters)} \\ y_{d=0} = \alpha_0 + \boldsymbol{z_M}\boldsymbol{\beta_{M0}} + \varepsilon_0 \text{ when } d^* \leq 0 \text{ (non} - \text{telecommuters)} \end{cases}$$

The two equations may be combined into the following "switching" model:

(5A) $y = (1-d)y_{d=0} + dy_{d=1}$, or

(5B) $y = \alpha_0 + d(\alpha_1 - \alpha_0) + \boldsymbol{z_M}\boldsymbol{\beta_{MO}} + d\boldsymbol{z_M}(\boldsymbol{\beta_{M1}} - \boldsymbol{\beta_{M0}}) + \varepsilon_0 + d(\varepsilon_1 - \varepsilon_0)$.

The assumption of homogeneity in self-selection bias may be equivalently expressed as the equivalency in the distribution of the errors ε_0 and ε_1, i.e. both $\mathrm{N}(0, \sigma^2)$, as well as both having the same correlation with υ. As a result, when there is homogeneity in unobserved tastes, Equation (5B) can be written as:

(6A) $y = \alpha_0 + d(\alpha_1 - \alpha_0) + \boldsymbol{z_M}\boldsymbol{\beta_{M0}} + d\boldsymbol{z_M}(\boldsymbol{\beta_{M1}} - \boldsymbol{\beta_{M0}}) + \varepsilon$,

and observable heterogeneity arises when $\boldsymbol{\beta_{M1}} \neq \boldsymbol{\beta_{M0}}$.

To ensure that the coefficient of *d* reflects the causal effect of *d*, we add and subtract $d\boldsymbol{\mu_{Z_M}}(\boldsymbol{\beta_{M1}} - \boldsymbol{\beta_{M0}})$, where $\boldsymbol{\mu_{Z_M}} = \boldsymbol{E}(\boldsymbol{z_M})$ is a vector of expected values of $\boldsymbol{z_M}$. Expressing $\boldsymbol{\delta}$=$(\boldsymbol{\beta_{M1}} - \boldsymbol{\beta_{M0}})$ and $\tau = \left[(\alpha_1 - \alpha_0) + \boldsymbol{\mu_{Z_M}}(\boldsymbol{\beta_{M1}} - \boldsymbol{\beta_{M0}})\right]$, we can re-write (6A) as:

(6B) $y = \alpha_0 + d\tau + \boldsymbol{z_M}\boldsymbol{\beta_{M0}} + d(\boldsymbol{z_M} - \boldsymbol{\mu_{Z_M}})\boldsymbol{\delta} + \varepsilon$.

In (6B), endogeneity arises due to the K+1 terms d and $d(\boldsymbol{z_M} - \boldsymbol{\mu_{Z_M}})$. To address endogeneity, we have to recognize that if $\boldsymbol{z_{IV}}$ is the instrument for *d*, then it is also a natural instrument for $d(\boldsymbol{z_M} - \boldsymbol{\mu_{Z_M}})$; hence estimated $\hat{d}_{i,P2}(\boldsymbol{z_{i,M}} - \bar{\boldsymbol{z}}_{\boldsymbol{M}})$ (where $\bar{\boldsymbol{z}}_{\boldsymbol{M}}$ is the sample estimate of $\boldsymbol{\mu_{Z_M}}$)

captures the good variation in $d(\boldsymbol{z_M} - \boldsymbol{\mu_{Z_M}})$ and may be substituted in (6B) to address endogeneity issues similar to the substitution of d by $\hat{d}_{i,P2}$.

Given the above, the H.2SPS method may be estimated using the stages listed earlier for 2SPS except for a modified third stage:

(i) Estimate $\hat{d}_{i,P1}$ using $\boldsymbol{z_{i,O}}$ and $\boldsymbol{z_{i,IV}}$

(ii) Run the following OLS regression and obtain $\hat{d}_{i,P2}$:

$$d_i \text{ on } 1, \hat{d}_{i,P1}, \boldsymbol{z_{i,M}}, \boldsymbol{z_{i,IV}}, \qquad i = 1, \dots, N.$$

(iii) Run the OLS regression:

$$y_i \text{ on } 1, \hat{d}_{i,P2}, \boldsymbol{z_{i,M}}, \hat{d}_{i,P2}(\boldsymbol{z_{i,M}} - \boldsymbol{\bar{z}_M}) \qquad i = 1, \dots, N.$$

The parameters summarizing treatment effects in (6B) are τ and $\boldsymbol{\delta}$, whose estimates are referred to as $\hat{\tau}_{H.2SPS}$ and $\boldsymbol{\hat{\delta}}_{H.2SPS}$.

In the Control Function approach (H.CF), the IMR $\lambda(\cdot)$ is substituted into (6B) in a process similar to that discussed in Section 4.3 to address endogeneity. The H.CF is a special case of the method detailed below, which assumes both observable and unobservable heterogeneity.

4.5 Heterogeneity in self-selection bias (Method: HH.CF)

In this section, we allow differences in self-selection bias between the two regimes, in addition to differences in observed tastes assumed in last section. This is equivalent to stating that the distributions of ε_1 & ε_0 , as well as the joint distributions of $(\varepsilon_1, \upsilon)$ and $(\varepsilon_0, \upsilon)$ are different. In addition, we continue to assume heterogeneity in observed tastes i.e. $\boldsymbol{\beta_{M1}} \neq \boldsymbol{\beta_{M0}}$. The error terms are assumed to be trivariate normally distributed as follows

$$(7) \qquad \begin{pmatrix} \varepsilon_1 \\ \varepsilon_0 \\ \upsilon \end{pmatrix} \sim N_3 \left[\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_1^2 & \sigma_{01} & \rho_1\sigma_1 \\ \sigma_{01} & \sigma_0^2 & \rho_0\sigma_0 \\ \rho_1\sigma_1 & \rho_0\sigma_0 & 1 \end{pmatrix} \right]$$

Making changes similar to those in (6A) & (6B) above, we can rewrite the switching equation (5D) as the following:

(8A) $y = \alpha_0 + d\tau + \boldsymbol{z_M}\boldsymbol{\beta_{MO}} + d(\boldsymbol{z_M} - \boldsymbol{\mu_{Z_M}})\boldsymbol{\delta} + \varepsilon_1 + (1-d)\varepsilon_0$

As previously, we add the control functions estimated from the IMR, to obtain:

(8B) $y = \alpha_0 + d\tau + \boldsymbol{z_M}\boldsymbol{\beta_{MO}} + d(\boldsymbol{z_M} - \boldsymbol{\mu_{Z_M}})\boldsymbol{\delta} + \rho_1\sigma_1 d[\phi(\boldsymbol{z\pi})/\Phi(\boldsymbol{z\pi})] +$

$$\rho_0\sigma_0(1-d)[\phi(\boldsymbol{z\pi})/\Phi(-\boldsymbol{z\pi})] + error$$

The method is estimated in the following steps:

(i) Estimate $\hat{d}_{i,P1}$ using (2B) for respondent $i, i = 1, \dots, N$.

(ii) Estimate the IMR terms,

$d_i\lambda(\boldsymbol{z_i\hat{\pi}})$ and $(1-d_i)\lambda(-\boldsymbol{z_i\hat{\pi}}), \qquad i = 1, \dots, N.$

(iii) Run the following OLS regression:

y_i on 1, d_i, $\boldsymbol{z_{i,M}}$, $d_i(\boldsymbol{z_{i,M}} - \boldsymbol{\bar{z}_M})$, $d_i\lambda(\boldsymbol{z_i\hat{\pi}})$, $(1-d_i)\lambda(-\boldsymbol{z_i\hat{\pi}})$

$i = 1, \dots, N.$

4.6 Treatment effects

We are interested in the population average effect of telecommuting adoption on travel behavior outcomes – the expected effect if a person is chosen randomly from the population and adopts telecommuting. In the causal inference literature, such an effect is referred to as the Average Treatment Effect (ATE), which is defined as:

$$\tau_{ATE} = E[y_{i,d=1} - y_{i,d=0}|\boldsymbol{z_{i,M}}],$$

where $y_{i,d=1}$ and $y_{i,d=0}$ denote the outcomes that will be realized by individual i if she adopts or does not adopt telecommuting respectively. A second estimand of interest is the Average Treatment Effect on the Treated (ATT or ATET), which reflects the treatment effect for a randomly chosen individual who has already adopted telecommuting:

$$\tau_{ATT} = E[y_{i,d=1} - y_{i,d=0}|\boldsymbol{z_{i,M}}, d = 1] .$$

Table 2. Estimation of treatment effects

#	**Method (relevant equations)**	**Average Treatment Effects (ATE)**	**Average Treatment Effects on the Treated (ATT)**
1	2SPS under homogeneity (Eq. 3)	$\hat{\tau}_{2SPS}$	Same as ATE
2	CF under homogeneity (Eq. 4, CF)	$\hat{\tau}_{CL}$	Same as ATE
3	2SPS under observable heterogeneity (Eq. 6, H.2SPS)	$\hat{\tau}_{H.2SPS} = (\hat{\alpha}_1 - \hat{\alpha}_0) + \bar{\boldsymbol{z}}_{\boldsymbol{M}}(\widehat{\boldsymbol{\beta}}_{\boldsymbol{M1}} - \widehat{\boldsymbol{\beta}}_{\boldsymbol{MO}})$	$ATE + \frac{1}{\sum_{i=1}^{N} d_i} \sum_{i=1}^{N} \left[d_i(\boldsymbol{z}_{\boldsymbol{i,M}} - \bar{\boldsymbol{z}}_{\boldsymbol{M}}) \widehat{\boldsymbol{\delta}}_{H.2SPS} \right]$
4	CF under observable heterogeneity (Eq. 6, H.CF)	$\hat{\tau}_{H.CF}$ (Same formula as $\hat{\tau}_{H.2SPS}$)	Same as above with $\widehat{\boldsymbol{\delta}}_{H.2SPS}$ replaced with $\widehat{\boldsymbol{\delta}}_{H.CF}$
5	CF under observable and unobservable heterogeneity (Eq. 8, HH.CF)	$\hat{\tau}_{HH.CF}$ (Same formula as $\hat{\tau}_{H.2SPS}$)	$ATE + \frac{1}{\sum_{i=1}^{N} d_i} \sum_{i=1}^{N} \left[d_i(\boldsymbol{z}_{\boldsymbol{i,M}} - \bar{\boldsymbol{z}}_{\boldsymbol{M}}) \widehat{\boldsymbol{\delta}}_{HH.CF} \right] + \frac{\hat{\rho}_1\hat{\sigma}_1 - \hat{\rho}_0\hat{\sigma}_0}{\sum_{i=1}^{N} d_i} \sum_{i=1}^{N} \left[d_i \left(\frac{\phi(\boldsymbol{z}\widehat{\boldsymbol{\pi}})}{\Phi(\boldsymbol{z}\widehat{\boldsymbol{\pi}})} \right) \right]$

Under assumptions of homogeneity between telecommuters and non-telecommuters, ATT is equal to ATE. If only the influence of observed covariates on outcome are assumed to differ between the two groups, then ATT differs from ATE by a function of the differences in the estimated magnitude of the influences $\boldsymbol{\beta}_{\boldsymbol{M1}}$ and $\boldsymbol{\beta}_{\boldsymbol{M0}}$. If the two groups differ also in terms of unobserved variables, then ATT differs from ATE by a function of the IMR capturing the nature of the self-selection bias.

Figure 2. Graphical representation of some of the key econometric concepts addressed in this study:

(A) Study motivation: Selection or omitted variable bias	(B) Validity of IVs: Traditional versus conditional instruments (Section 4.7)	(C) Bad conditioning: Potential for over-condition and endogenous selection bias (Section 4.8)
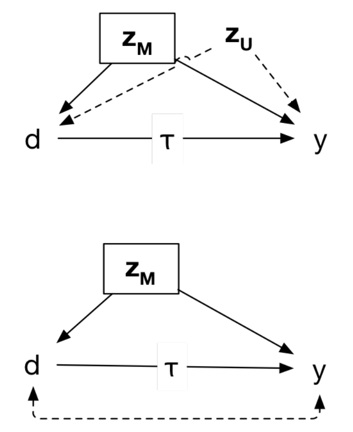	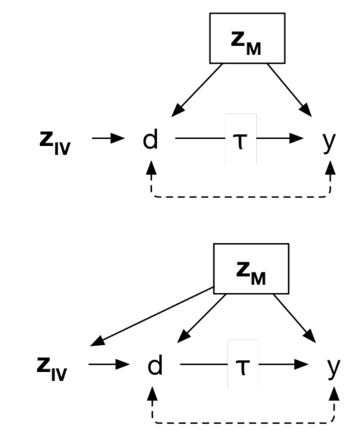	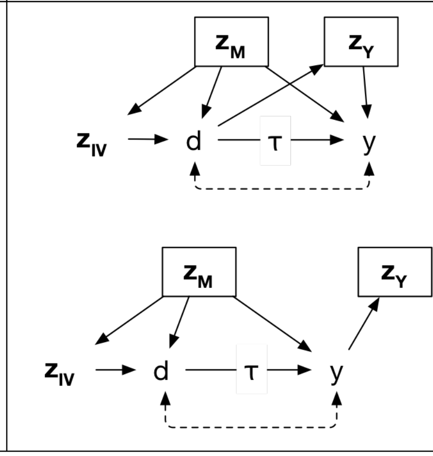

Notes: Uni-directional arcs represent causal effects, and bi-directional arcs represent correlations. A "boxed" variable implies conditioning, e.g. included explicitly as explanatory variable in a parametric analysis. (A) $\boldsymbol{z_U}$ influences both *d* and *y* (top). Because it is unobserved and cannot be conditioned, an estimate of effect τ will be conflated with bias (bottom). (B) A traditional instrument is exogenous to both *y* and $\boldsymbol{z_M}$ (top). If it is correlated with or influenced by $\boldsymbol{z_M}$, then conditioning on $\boldsymbol{z_M}$ renders it exogenous (bottom). The IV used in this study may be considered a "conditional" IV. (C) Conditioning on outcomes of *d* (top) or outcomes of *y* (bottom) leads to biased estimates of effect τ. Such outcomes should not be included as explanatory variables. A potential outcome to both *d* and *y* in this study is respondent's one-way commute distance.

For simplicity, we allowed only some of the components of $\boldsymbol{\beta_M}$ to differ between the telecommuting and non-telecommuting groups. Using an iterative process to identify statistically significant differences in estimates of $\boldsymbol{\beta_{M1}}$ and $\boldsymbol{\beta_{M0}}$, we finally adopted differences in coefficients of the socioeconomic group of the respondent.

4.7 Instrumental variables

Both the 2SPS and Control Functions methods rely on instrumental variables. As explained earlier in Section 4, an instrument must fulfill two key criteria. First, it should sufficiently explain the variation in *d*. We can explicitly test this criterion by checking the statistical significance of the estimated coefficient(s) of the IV(s) in a probit regression of *d* with the various exogenous confounders and all the instrumental variables. Additionally, we can check for change in pseudo R-squared with and without inclusion of the IV(s) while estimating the selection equation (probit regression). A sufficiently large change would imply that the IV(s) explain a large portion of the variation in *d* and are not *weak* instruments.

Second, the sole channel through which the IV(s) affect *y* is through the EEV (exclusion restriction). In other words, the IV should not affect *y* directly even though the empirical correlation between IV and *y* is unlikely to be zero because of the likely correlation of both *y* and IV with *d*. This assumption cannot be explicitly tested, but must be justified using substantive knowledge and a review of the literature in travel behavior. Recent literature has relaxed this assumption to suggest that if the instrument is only correlated with (or even influences) the travel behavior outcome through indirect channels, then conditional on the mediating variables, the instrument can be treated as a valid instrument.

We consider two potential IVs - availability of broadband internet at home, and potential to telecommute.

Zhu (2013); Zhu, Mason (2014) instrumented the telecommuting indicator with a variable indicating frequency of use of internet at home in their study of the effect of telecommuting on one-way commute distance and frequency of work and non-work trips based on the 2001 and 2009 NHTS. In our analysis, the coefficient of broadband availability is statistically significant thus technically meeting the first criterion for an IV. However, it is unclear whether the indicator can effectively discriminate between telecommuters and non-telecommuters, given the near ubiquity of internet access including broadband internet access. In our dataset of employed non-home-based workers, nearly 93% of respondents indicated availability of broadband at home. Among telecommuters, the ratio was 99%.

More importantly, we believe that the indicator may not meet the second criterion of a good IV. Mokhtarian (2009) identifies various ways in which internet availability and usage can facilitate both increases and decreases in travel, and not all these uses are mediated by telecommuting. For example, internet access provides an opportunity to substitute physical shopping with online shopping, thereby eliminating the need for some shopping trips. Given the existence of direct pathways through which internet access affects non-work travel, as well as the near ubiquity of access, we do not consider this variable a suitable IV.

The second potential IV is the potential to telecommute. Respondents were asked how much of their work can be theoretically undertaken from home assuming availability of necessary equipment and without considering employer policies regarding telecommuting. Options included "all", "most" "some", and "none". It can be easily argued that a person's assessment of the potential to undertake her work from home is unlikely to have a direct influence on travel behavior (second condition for IV). However, the potential to telecommute is not uncorrelated with other exogenous variables, and hence potentially affects travel indirectly (Figure 2B). Specifically, we find a strong correlation of the socioeconomic group on a person's assessment of the potential to undertake her work from home, thus violating the second condition for valid instruments. However, conditional on the exogenous variables $\boldsymbol{z_M}$, the variable may still be treated as a valid instrument; hence the term "conditional IVs" (Chen and Pearl 2014). We condition for these exogenous variables like the respondent's socioeconomic group by including them in the outcome model.

With respect to the first condition, we had to first address "separation" (Zorn 2005). This arises because the option of "none" in response to how of the work can be undertaken from home was available only to non-telecommuters; telecommuters could only choose from the three categories of "all", "most" and "some". In other words, knowing that a person cannot telecommute (aka "none"), leads to a perfect prediction that the person is a non-telecommuter. In presence of separation, the estimation of the selection model is either not possible or the estimates are invalid and unreliable. STATA®, the software package used in this study, removes the perfectly predicted observations while estimating the probit model on a standalone basis with *d* as the dependent variable. To address separation, we merged categories “some” and “none”. In the SI, we give the contingency table showing relationship between ability to telecommute and actual telecommuting status.

In the SI, we compare various goodness-of-fit measures of the two models with and without the proposed IV. The likelihood ratio chi-square increases by 327 for a change in 2 degrees of freedom (3 categories minus 1 reference category). The pseudo R-squared which also captures the model likelihood increases from 0.23 to 0.27. Other goodness-of-fit measures such as AIC and BIC also change substantially. Further, the estimated coefficients of the IV are large and statistically significant. As a result, we adopt the ability to telecommute as our instrumental variable.

4.8 Other potential sources of endogeneity

4.8.1: Reverse causality

Decisions concerning telecommuting adoption and various dimensions of travel behavior such as primary commute mode, household vehicle holdings, and overall annual travel are dynamic processes that play out over multiple time periods. This creates the potential for a reverse causality bias when the data for analysis is cross-sectional, capturing telecommuting and travel behavior patterns as a snapshot in time. Specifically, high levels of travel even before adoption of telecommuting may encourage the adoption of telecommuting so as to optimize one’s overall travel budget. This influence runs in a direction opposite to the effect of interest – that of telecommuting on the level of travel – and thus, if only the latter direction is modeled, the estimated effect of telecommuting on travel is apt to be biased upward.

The assertion of existence of this two-way relationship is supported by comparing the average frequency of business trips and level of business travel between the two groups. As mentioned before, 81% of the non-telecommuters in our sample did not report any business trips in their one-week travel diary, while the corresponding number for telecommuters is 55%.

Among those who reported business trips, telecommuters undertook around 50% more business travel than non-telecommuters (142 versus 92 miles/week). We treat business travel as an exogenous covariate ($\boldsymbol{z_M}$) thereby implicitly making the assumption that such travel influences adoption of telecommuting as well as total non-work and commute travel. We believe that doing so should greatly reduce a reverse causality bias.

In addition to business travel, higher levels of personal and commute travel may also influence the propensity to telecommute, but given the purpose of this study, these clearly cannot be treated as exogenous. A natural idea, then, would be to use a Structural Equation Modeling (SEM) framework, specifying a non-recursive or reciprocal model with uni-directional arcs representing causal effects from d to y, and from y to d; and estimating the two equations simultaneously. However, non-recursive models have one key limitation: because the endogenous variable d is binary, the equation for y cannot have d (the actual treatment status) as an explanatory variable, but only d^*, which is the underlying latent continuous variable representing the propensity to enroll in the treatment. (In brief, the reason for this is that the alternative specification leads to a logical inconsistency if we were to insert the equation for y into the equation for d^* and simplify, because it leads to an equation in which the variable d – which can only take on the values 0 and 1 – can be directly expressed as a function of its corresponding latent variable d^* – which is continuous-valued). In other words, we cannot directly estimate the effects of telecommuting adoption itself (as is conceptually more appropriate) but rather the effects of the *propensity to adopt* telecommuting (Mishra et al. 2017; Madalla 1983)

In view of the serious conceptual deficiency of such a model, and in view of controlling for the level of business travel, which is likely to be a major source of reverse causality, we elected not to adopt an SEM in this context. Nevertheless, reverse causality remains, to some extent, a confounding factor in our results. If our results had shown telecommuting to be increasing travel, we would have said that reverse causality may be inflating that estimated impact (i.e. both y → d and d → y would have the same sign). Since our best results (see Section 5) show that telecommuting *decreases* travel, however, the likely impact of reverse causality is to *dampen* our estimate of that increase (the unmodeled y → d effect is positive, partly counteracting the modeled negative d → y effect). In other words, if we were able to properly account for reverse causality, we would likely see an even stronger substitution effect of telecommuting on travel.

4.8.2 Over-conditioning

Conditioning on a variable that is an outcome of either d or y, as indicated in Figure (2C) will lead to "bad control" and bias estimates of τ (Angrist and Pischke 2008). Elwert, Winship (2014) call the former "over control bias" and the latter "endogenous selection bias". To avoid such over-conditioning, we need to judiciously exclude variables that may induce such bias.

In our list of available variables (Table 1), one-way commute distance presents such a risk. We initially included the variable because a number of papers explaining adoption and frequency of telecommuting have modeled one-way commute distance as an explanatory variable. Further, in terms of time scale, choices about residential location are longer-term decisions compared to adoption of telecommuting, thereby justifying the assumption of exogeneity. The above reasoning necessitates inclusion of the variable as a component of $\boldsymbol{z_M}$ to control for observed differences in traits or confounder bias. However, a few papers have found a causal effect of telecommuting on one-way commute distance. For example, Zhu (2013) concludes that telecommuting leads to households choosing residential locations farther away from the workplace. Analyzing the temporal order of telecommuting engagement and residential relocation, (Ory and Mokhtarian 2006) show that those who are telecommuting and then move actually tend to relocate closer to their workplace, whereas those who begin telecommuting following a residential relocation tended to have moved much farther from their workplace.

It may be argued that a misguided exclusion of an observed confounder to avoid bad control should not exacerbate the endogeneity problem if steps are being taken to address omitted variable bias. In other words, the excluded variable may be considered part of $\boldsymbol{z_U}$, whose omission will be addressed using the methods discussed above. Given the above, we exclude the one-way commute distance variable as a component of $\boldsymbol{z_M}$ in our models.

5. Results

Table 3 compares the estimates of average treatment effects (ATE) of telecommuting on annual total travel (not including business trips), commute-only travel, and finally non-work travel, by the five methods summarized in Table 2. We also present the estimates provided by a regression model, which includes all components of $\boldsymbol{z} = [\boldsymbol{z_M}, \boldsymbol{z_{IV}}]$ as explanatory variables. The regression model only controls for observed differences in traits and tastes of telecommuters and non-telecommuters, but does not attempt to control for any unobserved differences. Similar to 2SPS and CF, we allowed only some of the components of $\boldsymbol{\beta_M}$ to differ between the telecommuting and non-telecommuting groups – specifically, those pertaining to the socioeconomic group of the respondent.

The detailed results, including estimates of coefficients of $\boldsymbol{z_M}$ and $\boldsymbol{z_{IV}}$ are available in the SI.

Table 3. (A, top): Average percentage change in annual non-work and commute distance travelled with adoption of telecommuting; (B, bottom): Estimates of self-selection bias

		All trip purposes (except business)		Commute trips	Non-work trips
		All modes	**Cars & van**	**Cars & van**	**Cars & van**
		$y1$	$y1_{car}$	$y2_{car}$	$y3_{car}$
	Regression	1% (-6% to 9%)	1% (-6% to 8%)	-11%* (-18% to -3%)	5% (-8% to 19%)
1	2SPS	-19%* (-28% to -9%)	-15%* (-26% to -2%)	-23%* (-34% to -10%)	-4% (-21% to 16%)
2	CF	-15%* (-24% to -6%)	-12%* (-21% to -2%)	-25%* (-34% to -14%)	-12% (-31% to 12%)
3	H.2SPS	-13%* (-24% to -1%)	-13%* (-24% to 0%)	-33%* (-55% to -1%)	-9% (-29% to 15%)
4	H.CF	-24%* (-35% to -11%)	-19%* (-31% to -5%)	-32%* (-45% to -17%)	-13% (-35% to 17%)
5	HH.CF	-26%* (-37% to -13%)	-20% (-33% to 6%)	-33%* (-45% to -18%)	-14% (-37% to 17%)

	Method	All trip purposes (except business)		Commute trips	Non-work trips
		All modes	**Cars & van**	**Cars & van**	**Cars & van**
2	CF	$\hat{\rho} = 0.17$* (0.07 to 0.26)	$\hat{\rho} = 0.12$* (0.02 to 0.22)	$\hat{\rho} = 0.14$* (0.04 to 0.24)	$\hat{\rho} = 0.05$ (-0.05 to 0.15)
4	H.CF	$\hat{\rho} = 0.23$* (0.11 to 0.34)	$\hat{\rho} = 0.17$* (0.05 to 0.28)	$\hat{\rho} = 0.17$* (0.05 to 0.29)	$\hat{\rho} = 0.08$ (-0.04 to 0.20)
5	HH.CF	$\hat{\rho}_0 = 0.11$ (-0.06 to 0.28) $\hat{\rho}_1 = 0.28$* (0.15 to 0.40)	$\hat{\rho}_0 = 0.11$ (-0.05 to 0.27) $\hat{\rho}_1 = 0.19$* (0.06 to 0.32)	$\hat{\rho}_0 = 0.12$ (-0.06 to 0.29) $\hat{\rho}_1 = 0.19$* (0.06 to 0.31)	$\hat{\rho}_0 = 0.06$ (-0.11 to 0.24) $\hat{\rho}_1 = 0.09$ (-0.05 to 0.22)

Note: Figures in parentheses reflect 95% CIs. "*" represents estimates which are statistically significant at $\alpha = 0.05$ or lower. Estimates of $y2$ and $y3$ are similar to $y2_{car}$ and $y3_{car}$ respectively, and not summarized above for simplicity. Population: All full-time employed respondents with fixed work location outside home. Effect estimates may be considered as average treatment effect or ATE.

The regression estimates indicates that telecommuting adoption has a substitutive effect on commute travel – the mean estimate is quite large (almost 11% reduction in total annual travel) and statistically significant. The estimated effects on total and non-work travel are positive (complementarity); however, the mean estimates are close to zero and not statistically significant. In comparison, Figure 2 straightforwardly indicates that telecommuters travel substantially more than non-telecommuters – for both commute and non-work purposes. Hence, it may be concluded that that some of the observed differences in travel between the two groups may be explained by observed differences in characteristics and tastes.

Methods that seek to control for unobserved differences in characteristics estimate a substitution effect on commute travel that is larger in magnitude compared to regression estimates - a reduction in annual commute travel by car or van ($\boldsymbol{y2_{car}}$) by one-quarter to one-third. The effect estimates for non-work travel are centered below zero indicating a substitution effect in contrast to regression (non-work) estimates; however, the estimates are not statistically significant. Further, the estimates have large standard errors. For example, the 95% confidence interval for estimates of reduction in $\boldsymbol{y3_{car}}$ by our best model HH.CF ranges from -37% to +17%. This indicates large variability in the effect of telecommuting adoption in the population. Further, it also suggests that both complementary and substitutive effects are in play, with the relative proportions of these effects varying among people.

The nature and magnitude of selection bias is also evident from the estimates of ρ as summarized in Table 3B. Rho estimates for all methods and for all travel variables are positive, indicating that unobserved traits and tastes influence the adoption of telecommuting in the same direction as they influence travel volumes. Further, the estimates are either significant (at $\alpha = 0.05$) or "marginally" significant (given that the lower bound of the 95% CI is close to zero). This, at least partially, explains the observed higher travel volumes of telecommuters compared to non-telecommuters (Figure 1), the more negative effect estimates by the CF or 2SPS methods compared to regression (Table 3A), and the complementarity effect sometimes concluded in the literature on telecommuting.

In Figure 3, we present the results of just our "best" model (HH.CF) and the regression model, disaggregated by trip purpose and mode (car/van or all modes). Further, we disaggregate the results based on current telecommuting status to identify differences in the average treatment effect on the treated and on the controls (ATT versus ATC or ATUT). The ATT may be considered as the reduction in travel compared to a counterfactual scenario where telecommuters do not adopt telecommuting. On the other hand, the ATC may be considered as the potential reduction in travel if non-telecommuters were to adopt telecommuting. The average treatment effect or ATE is the weighted average of these two effects – where the weights are the proportions of the two groups in the sample (or population).

Figure 3A reports the ATT and ATC in terms of percentage reduction in distance traveled for a given category, which (in view of the log transformation of the distance dependent variable) is the direct way the effects are computed. But it is also of interest to evaluate the effects in terms of distance traveled per se, since ultimately it is that measure which relates directly to fuel consumption, emissions, and congestion. Averaging individual-specific percentage reductions can be deceiving, since changes from a small baseline amount of travel can appear quite large in percentage terms but be negligible in terms of system impacts, and conversely. Accordingly, Figure 3B reports the two effects in terms of distance traveled itself. Table 4 summarizes the mean and median effects.

Figure 3. Estimated effect of telecommuting adoption on annual travel using two methodologies – disaggregated by trip purpose, travel modes, and current telecommuting status. Effect estimates in terms of percentage reduction (Top) and absolute miles reduction (Bottom) in annual travel

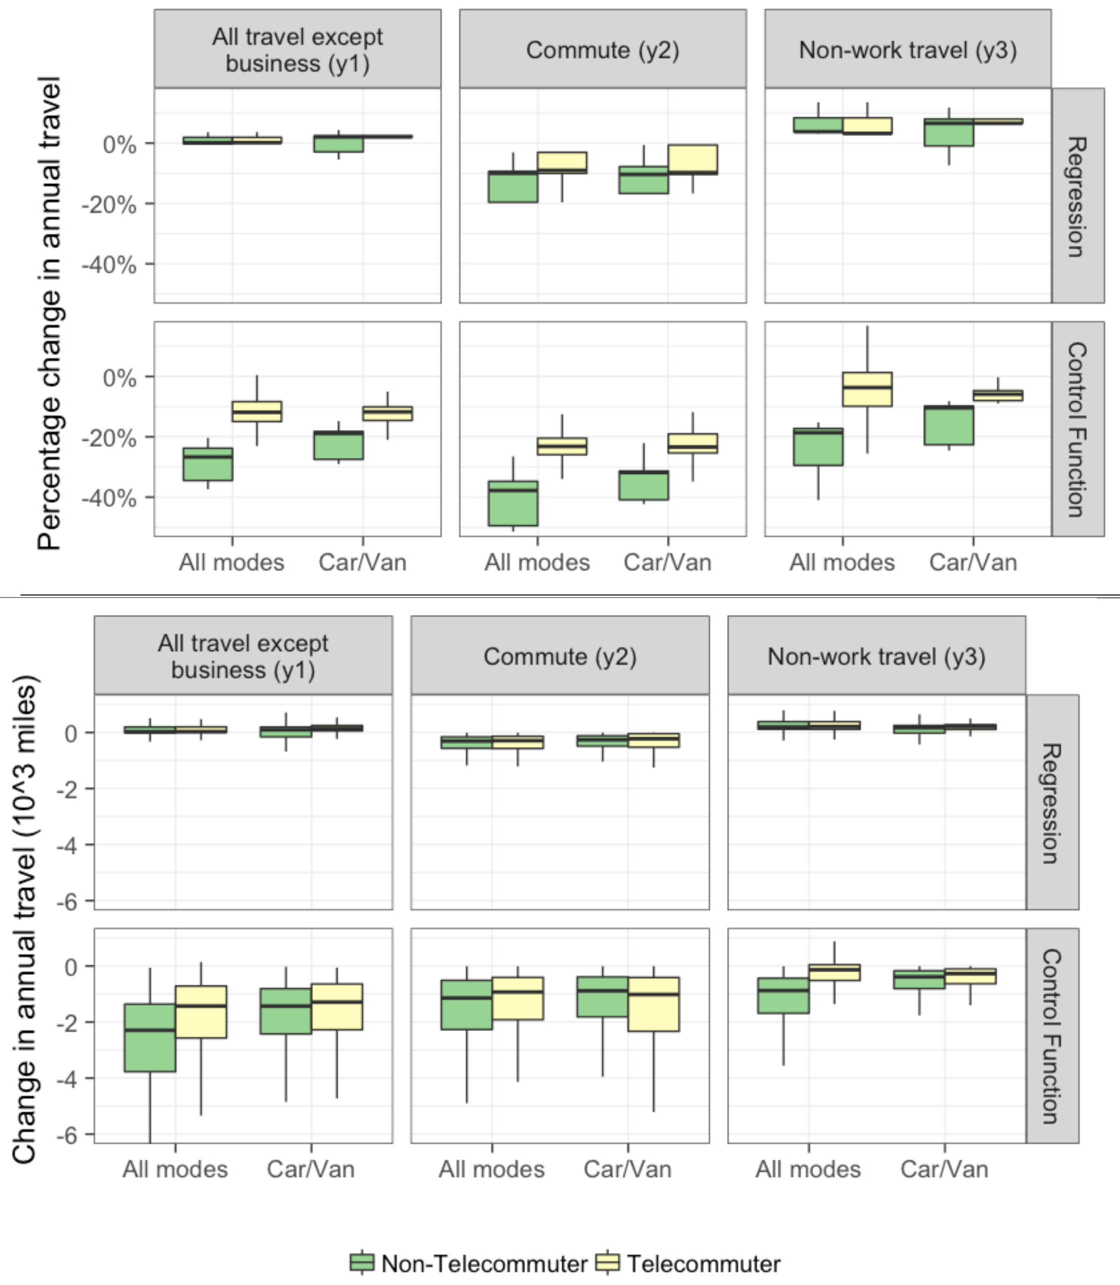

Note: Effect estimates are presented for our "best" model (HH.CF) which controls for observed and unobserved differences in both traits and tastes. We present the estimates of the regression model, which only controls for observed differences in traits and tastes, to highligh the nature and extent of selection bias. The effect estimates for telecommuters may be considered as average treatment effect on the treated (ATT); the estimates for non-telecommuters as the average treatment effect on control (ATC) or untreated (ATUT). Also, "All travel (y1)" excludes business related travel.

As a result of the adoption of telecommuting, annual total (commute plus non-work) travel by car/van reduced by 12 (~1,300 miles) for telecommuters and could potentially fall by 22% (~1,850 miles) for non-telecommuters. The difference in estimated effect between the two groups is wider for total travel by all modes: 12% (~1,900 miles) and 28% (~2,300 miles) for telecommuters and non-telecommuters respectively. The effects summarized above are mean estimates. Commute travel accounts for 60-70% of the above potential reduction in total annual travel for non-telecommuters, but almost 70-80% of the observed reduction for telecommuters. This may be partly explained by the much larger one-way commute distance of the latter (average of 20 miles compared to 10 miles for non-telecommuters). The mean annual reduction in non-work travel by car is around 530 and 655 miles for telecommuters and non-telecommuters respectively.

The small effect size along with high standard errors summarized in Table 3A leads us to believe that telecommuting effects on non-work trips are simultaneously substitutive and complementary. The theoretical justifications for both types of effects with respect to non-work travel have been discussed extensively in the literature and summarized well by (Kim, Choo and Mokhtarian 2015). For example, income effects or resource savings may promote complementarity, in that the savings in travel time and costs enabled by telecommuting can be converted into expenditures for other travel purposes. Additionally, telecommuters' reduced opportunities for social interaction at work may incentivize increased participation in community activities and civic engagement, thus generating new travel. Substitution effects, on the other hand, may result from a shrinking of the activity space and concentration of destinations around the residence as individuals spend more time based at home instead of the workplace. Telecommuters may thus discover new destinations for daily activities closer to home, thereby reducing non-work travel. It is unlikely that only one of these effects are in play for all individuals; we believe that both effects are working simultaneously and probably largely negating each other. The strength of the opposing effects may vary across individuals, leading to the observed high standard errors.

Table 4. Mean and median estimates of the effect of telecommuting adoption on annualized travel based on our "best" model (HH.CF):

		Tele-commutes?	Effect (change in miles)		Effect (percent change)	
			Mean	**Median**	**Mean**	**Median**
All Travel	All modes	No	-2,912	-2,293	-28%	-27%
(except business		Yes	-1,905	-1,428	-12%	-12%
travel)		Combined	-2,797	-2,179	-26%	-26%
	Car	No	-1,861	-1,431	-22%	-19%
		Yes	-1,698	-1,288	-12%	-12%
		Combined	-1,842	-1,417	-21%	-19%
Commute	All modes	No	-1,766	-1,141	-40%	-38%
travel		Yes	-1,467	-929	-23%	-23%
		Combined	-1,732	-1,117	-38%	-36%
	Car	No	-1,409	-887	-34%	-32%
		Yes	-1,743	-1,016	-23%	-23%
		Combined	-1,447	-900	-33%	-32%
Non-work	All modes	No	-1,322	-876	-22%	-19%
travel		Yes	-355	-134	-5%	-4%
		Combined	-1,212	-791	-20%	-18%
	Car	No	-655	-386	-15%	-10%
		Yes	-530	-273	-8%	-6%
		Combined	-641	-374	-14%	-10%

6. Conclusions

This study sought to examine how the adoption of telecommuting influences annual non-commute and commute travel using a large household travel survey. The study controls for observed and unobserved differences in traits and tastes between telecommuters and non-telecommuters by using two parametric approaches that rely on instrumental variables – the Two-Stage Predictor Substitution (2SPS) and the Control Function (CF) methods. Both approaches indicate that telecommuting has a substitution effect and leads to lower commute and non-work related travel. The CF method also allows us to identify the nature and magnitude of self-selection bias – the net effect of unmeasured factors influences both the adoption of telecommuting and travel volumes in the same direction. The strength of this correlation is somewhat stronger for telecommuters than non-telecommuters.

The finding of a substitution effect of telecommuting contrasts with the directly-observed differences in travel volumes between telecommuters and non-telecommuters. As we show in Figure 1, telecommuters travel substantially more than non-telecommuters for both commute and non-work related travel. Controlling for only the observed differences in traits and tastes between the two groups reduces the difference to close to zero (Figure 3, Regression results). Only after we further control for unobserved differences in characteristics (aka self-selection bias), can we identify the substitution effects of telecommuting.

There are several limitations of this study. First, it may be argued that a more appropriate treatment indicator may be the frequency of commuting to one's workplace location rather than the adoption of telecommuting per se. A non-telecommuting respondent working on a 10-hour/day 4-day work schedule commutes as frequently as a telecommuter who works on an 8-hour/day schedule and telecommutes once a week. However, no data is available about a respondent's typical work schedule and average frequency of commuting to the workplace. We suggest that future surveys should also measure the work schedule of all employed respondents in addition to their telecommuting status. Second, our treatment indicator does not consider differences in the frequency of telecommuting by telecommuters, and adopts a simple dichotomous yes/no indicator. It is reasonable to assume the reduction in travel volumes to be more closely associated with the frequency of telecommuting than with the fact of telecommuting at all. In a medical context, a similar example is using level of drug dosage instead of an indicator that only captures whether the drug was given or not.

Third, we do not fully control for potential endogeneity due to reverse causality, which arises when an "explanatory" variable is simultaneously a function of the "dependent" variable it is supposed to explain – that is, when one variable is both a cause and an effect of another. In addition to telecommuting influencing the level of travel (the effect of interest in this study), high levels of travel even before adoption of telecommuting may influence the adoption of telecommuting. We partly control for this bias by conditioning on frequency of business trips given the observation in this dataset that telecommuters are more likely to report business trips in their travel diary (40% versus 10% of respondents) and also report longer travel distances. However, this does not control for the potential influence of greater *personal* travel on the propensity to adopt telecommuting. Future research should address the above limitations for consistent and unbiased estimates of telecommuting. In Section 4.8, we argued that had we properly accounted for reverse causality, we would have likely seen an even stronger substitution effect of telecommuting on travel.

The study demonstrates an approach to drawing causal inferences from publicly available travel surveys like the UK-NTS and NHTS. These surveys are observational and cross-sectional in nature, leading to the potential for various biases while drawing causal inferences. Further, these surveys collect little or no information on attitudes and preferences towards such issues as the environment, driving, transit use, and urban living, although such attitudes have been shown to considerably improve the ability to explain people's travel and residential location decisions (Handy, Cao and Mokhtarian 2005). These limitations may partly explain why most studies on the causal effects of telecommuting have been restricted to studies with purpose-built survey instruments with retrospective reporting. Such studies have small sample sizes with limited coverage across geographies or other demographic dimensions, a declared focus on telecommuting evaluation that may bias responses, and sometimes-proprietary datasets. We hope that our research will encourage more accurate evaluations of causal effects from publicly available national and regional travel surveys with large sample sizes and wide geographic and demographic scopes.

7. Acknowledgements

We would like to thank Kathryn Simpson at the School of Social Sciences University of Manchester and various statisticians at the UK Department for Transport for answering our questions pertaining to use of the UK NTS dataset. The first author was encouraged to pursue this study after extensive discussions with Kenneth Laberteaux from the Toyota Research Institute of North America. The views and opinions expressed in this paper are those of the authors alone.

The authors declare no conflict of interest

8. References

Andreev, P., Salomon, I., Pliskin, N.: Review: State of teleactivities. Transportation Research Part C: Emerging Technologies **18**(1), 3-20 (2010)

Angrist, J., Pischke, J.-S.: Mostly harmless econometrics: An empiricist's companion. Princeton University Press, (2008)

Bushway, S., Johnson, B.D., Slocum, L.A.: Is the magic still there? The use of the Heckman Two-Step Correction for selection bias in Criminology. Journal of Quantitative Criminology **23**(2), 151-178 (2007)

Cerulli, G.: Econometric evaluation of socio-economic programs: Theory and applications, vol. 49. Springer-Verlag Berlin Heidelberg, (2015)

Chen, B., Pearl, J.: Graphical tools for linear structural equation modeling. Psychometrika (2014)

Elwert, F., Winship, C.: Endogenous Selection Bias: The problem of conditioning on a collider variable. Annual Review of Sociology **40**, 31-53 (2014). doi:10.1146/annurev-soc-071913-043455

Handy, S., Cao, X., Mokhtarian, P.: Correlation or causality between the built environment and travel behavior? Evidence from Northern California. Transportation Research Part D: Transport and Environment **10**(6), 427-444 (2005). doi:10.1016/j.trd.2005.05.002

He, S., Hu, L.: Telecommuting, income, and out-of-home activities. Travel Behaviour and Society **2**(3), 131-147 (2015)

Kim, S.-N., Choo, S., Mokhtarian, P.L.: Home-based telecommuting and intra-household interactions in work and non-work travel: A seemingly unrelated censored regression approach. Transportation Research Part A: Policy and Practice **80**, 197-214 (2015)

Madalla, G.: Limited-dependent and qualitative data in econometrics. vol. 3. Cambridge University Press, (1983)

Mishra, G.S., Mokhtarian, P.L., Clewlow, R.R., Widaman, K.F.: Addressing the joint occurrence of self-selection and simultaneity biases in the estimation of program effects based on cross-sectional observational surveys - Case study of travel behavior effects in carsharing. In. Manuscript submitted for publication, (2017)

Mokhtarian, P.: If telecommunication is such a good substitute for travel, why does congestion continue to get worse? Transportation Letters **1**(1), 1-17 (2009)

Mokhtarian, P., Handy, S.L., Salomon, I.: Methodological issues in the estimation of the travel, energy, and air quality impacts of telecommuting. Transportation Research Part A: Policy and Practice **29**(4), 283-302 (1995)

Mokhtarian, P., Salomon, I.: Modeling the desire to telecommute: The importance of attitudinal factors in behavioral models. Transportation Research Part A: Policy and Practice **31**(1), 35-50 (1997)

Mokhtarian, P.L., Cao, X.: Examining the Impacts of Residential Self-selection on Travel Behavior: A Focus on Methodologies. Transportation Research Part B: Methodological **42**(3), 204-228 (2008). doi:DOI: 10.1016/j.trb.2007.07.006

Mokhtarian, P.L., van Herick, D.: Quantifying residential Self-Selection effects: A review of methods and findings from applications of Propensity Score and Sample Selection approaches. Journal of Transport and Land Use **9**(1) (2016)

Ory, D.T., Mokhtarian, P.L.: Which came first, the telecommuting or the residential relocation? An empirical analysis of causality. Urban Geography **27**(7), 590-609 (2006)

Popuri, Y.D., Bhat, C.R.: On modeling the choice and frequency of home-based telecommuting. (2003)

Rose, D., Pevalin, D.J., O'Reilly, K.: The National Statistics Socio-economic Classification: origins, development and use. Palgrave Macmillan Basingstoke, (2005)

Singh, P., Paleti, R., Jenkins, S., Bhat, C.R.: On Modeling Telecommuting Behavior: Option, Choice, and Frequency. Transportation **40**(2), 373-396 (2013)

StataCorp: Stata Treatment Effects reference manual - Release 14. StataCorp College Station, Texas (2015)

Taylor, E., Humphrey, A., Pickering, K., Tipping, S.: UK National Travel Survey 2012 - Technical Report. In. NatCen Social Research, Prepared for The Department for Transport, UK, London, (2013)

Toomet, O., Henningsen, A.: Sample selection models in R: Package sampleSelection. Journal of statistical software **27**(7), 1-23 (2008)

UK Department for Transport: National Travel Survey trip chaining: 2002-2014. In. London, (2015a)

UK Department for Transport: UK National Travel Survey data extract user guide: 1995-2014. In. London, (2015b)

Walls, M., Safirova, E., Jiang, Y.: What drives telecommuting? Relative impact of worker demographics, employer characteristics, and job types. Transportation Research Record: Journal of the Transportation Research Board **2010**(1), 111-120 (2007)

Wooldridge, J.: Econometric analysis of cross section and panel data. MIT Press, Cambridge, MA (2010)

Wooldridge, J.: Control function methods in applied econometrics. Journal of Human Resources **50**(2), 420-445 (2015)

Zhu, P.: Telecommuting, household commute and location choice. Urban Studies **50**(12), 2441-2459 (2013)

Zhu, P., Mason, S.G.: The impact of telecommuting on personal vehicle usage and environmental sustainability. International Journal of Environmental Science and Technology **11**(8), 2185-2200 (2014)

Zorn, C.: A solution to separation in binary response models. Political Analysis **13**(2), 157-170 (2005)