Climate Mitigation and Spatial Distribution of Automobile Demand: The Role of Income, Public Transit, and Portfolio Preferences

Yoshifumi Konishi Department of Economics Keio University

Sho Kuroda^{*} Graduate School of Systems and Information Engineering University of Tsukuba

Shunsuke Managi Department of Urban and Environmental Engineering School of Engineering, Kyushu University

Version: December 2021

^{*}Corresponding author. Mailing address: 1-1-1 Tennodai, Tsukuba, Ibaraki 305-8573, Japan. E-mail: s1530138@u.tsukuba.ac.jp.

Acknowledgement: Our special thanks goes to Shin Hong Joo and Akari Ono, who provided valuable assistance in preparing the dataset for this study. The study was in part supported by the financial support from Japan Society for the Promotion of Science (JSPS), Grant-in-aid for Scientific Research (*Grant numbers: 18K01562* and 21H00712). We also thank Takuya Hasebe, Taiju Kitano, Toshifumi Kuroda, Ryo Nakajima, and other seminar participants at the Japanese Economic Association 2020 Annual Meeting, Sophia University's Economics Seminar, the 14th Empirical Moral Science Workshop, and the 35th Applied Regional Science Conference, for their valuable comments.

Abstract: We empirically characterize how automobile demand varies over geographic space and how it affects the economic consequences of climate mitigation policies. We augment a discrete-continuous choice model in ways that account for geographic distribution of incomes, public transit, and portfolio preferences, and show that our model outperforms a *naive* random-coefficient model in explaining demand *stickiness* over geographic space. In particular, the model allows us to resolve two empirical puzzles in Japan: Overall price elasticity of demand for vehicle ownership increases with vehicle size; invariance of demand for hybrid vehicles with respect to public transit density. The estimated model substantiates the importance of this spatial demand heterogeneity for policy evaluation: Carbon tax has a larger CO_2 -reducing impact in non-urban settings, yet the effect flattens out as the transit density further declines as the vehicle demand becomes increasingly *sticky*; Consequently, the welfare loss from carbon tax is greatest in the lowest density areas; Eco-car sharing can mitigate this welfare penalty while remaining equally effective; Feebates perform poorly relative to either policy.

JEL Codes: H23, H31, L62, Q54

Key Words: Automobiles, carbon emissions, carbon tax, control function, discrete-continuous choice, energy efficiency, feebates, hybrid vehicle, portfolio effect, public transportation

1. Introduction

Carbon dioxides (CO₂) emissions pose a significant risk of climate change. Road transportation accounts for a large share of the global CO₂ emissions today. In theory, a carbon tax on gasoline consumption can fully restore economic efficiency since it forces firms and consumers to internalize the social cost of climate change in all economic margins: from residential/employment choice to transport mode choice, and to vehicle ownership and utilization. Economists have been increasingly aware, however, of the difficulties associated with reducing carbon emissions from road transportation [see Anderson *et al.* (2011) or Knittel (2012) for a review of the issues]. In the U.S., for example, vehicle miles travelled nearly doubled from 1970 to 2009 while the fleet average fuel economy (of only new cars) improved only at a moderate rate over the same period (Knittel, 2012). Associated with this increase in vehicle miles is the rapid urban sprawl: i.e., urban development in geographically sparse, low density areas (Glaeser and Kahn, 2010). The form of cities, roads, and public transit networks we observe today is the result of this development, so is the demand for vehicle transport [e.g., Beaudoin and Lawell (2018); Bento *et al.* (2005); Boarnet and Crane (2001); Duranton and Turner (2011)].

The theme of this manuscript is an important question that arises from this observation: How does automobile demand vary over geographic space? Addressing this question empirically is important, at least on three accounts. First, the question is of economic significance by itself because we know relatively little about the geographic heterogeneity of automobile demand, despite that there is a large literature on economic studies of automobile demand. We are, for example, not sure how demand elasticities differ between urban, suburban, and rural contexts, for what reasons — do they differ because of income, landscape, public transit, or other systematic differences in preference structures? Second, it matters for economic efficiency because the economic impact of a (price-based) climate mitigation policy can be highly spatially heterogeneous precisely due to this demand heterogeneity over space. The key is the joint distribution of incomes, public transit networks, and preference structures and their interactions over geographic space. For example, demand for vehicle ownership and utilization is low in urban areas where a dense public transit network is available. This makes automobile demand more price elastic, *ceteris paribus*. On the other hand, incomes tend to be high in urban areas where abundant employment opportunities exist (even after adjusting for the cost of living).¹ This makes the demand less price elastic. Thus, the interaction between the two effects alone tends to generate substantial heterogeneity in demand elasticity across geographic areas, so is the effect of a climate mitigation policy. Third, it is

¹There is a large literature showing that real wages are not equalized across regions.

important for equity reasons. Over the last few decades, there has been a growing interest among environmental economists in the equity implications of climate mitigation policies [e.g., Dorband *et al.* (2019); Fremstad and Paul (2019); Goulder *et al.* (2019); Grainger and Kolstad (2010); Maestre-Andres *et al.*, (2019)]. Most of the studies, however, examine the regressivity of climate policies and do not pay sufficient attention to the *joint* distribution of incomes, public transit, and preference structures over geographic space. For example, efforts to reduce carbon emissions by forming a smart/compact city may raise the cost of living in urban areas, forcing some of lower income populations to live in non-urban areas. Even a moderate carbon tax may hurt these disadvantaged subpopulations severely, not only because transportation costs represent a larger share of their incomes but also because they may have no other means to travel.

With this general theme in mind, this manuscript estimates a spatially explicit model of automobile demand in the discrete-continuous choice framework à la Dubin and McFadden (1984), using spatially rich survey data in Japan. As in Train (1986), our model accounts for a sequence of three choices: the number of vehicles owned, the class/type of each vehicle owned, and the vehicle kilometers traveled (VKT) for each vehicle owned. What is new here is that we incorporate three new aspects into the model and do so in a manner that is theoretically consistent with both the theory of travel demand (Domencich and McFadden, 1975) and the error-component formulation of mixed logit (Brownstone and Train, 1998). First, we introduce *portfolio considerations* in a manner analogous to Gentzkow (2007) and Wakamori (2015). That is, we explicitly model the correlation between choices of the first and the second cars by adding the terms that capture utility from having a particular combination of vehicles. Adding the portfolio effect allows us to model intricate behavioral responses that seem quite important in our empirical context, which we shall turn to below. Second, following the spirit of random-coefficient logit, we allow the parameters on (indirect) utility to depend explicitly on a measure of public transit density.² This formulation generates realistic substitution patterns that are explicitly linked to public transit. For example, a consumer who has a high valuation of fuel economy because her access to public transit is limited is allowed to substitute to a less expensive alternative, such as a keicar, that has less but similar fuel economy when the price a hybrid car is too high. Third, the resulting model produces the high dimensional sample correction terms that enter the vehicle utilization equation. To address it, we employ Dahl's control function approach (2002). By this, we account for two types of correlation in the VKT equation due to unobservables: one between the uses of multiple cars and the other between car ownership and utilization. We estimate this model, using a large, nation-wide internet survey we conducted in 2016. The survey contains a

²We explain how we construct this measure in **Section 3**.

usable sample of approximately 100,000 households, and hence, we have a sufficiently large subsample within each decile of public transit density. The covariates' variations that come with it allows us to estimate the demand parameters that vary by that density.³

Our modeling strategy is motivated, not only by the key theme of the manuscript discussed above, but also by our desire to resolve two empirical puzzles we observe in Japan. First, Konishi and Zhao (2017) find that in Japan, the price elasticities of automobile demand are smaller for small-sized vehicles than for large-sized vehicles. This is puzzling in that low-income households tend to buy smaller cars, and hence, the demand for these cars is expected to be price-elastic, in principle. Second, in our data (to be described below), we observe that the demand for hybrid cars is virtually invariant with respect to the public transit density. This is also puzzling because the demand for vehicle transport increases quite sharply, so does the demand for fuel economy, as public transit becomes more sparsely available. As it turns out, the key to resolving both of these puzzles lies with the joint distribution of income, public transit, and portfolio preferences over geographic space. Households in low density areas tend to buy keicars⁴, which are cheap and fuel efficient, instead of hybrid cars, which are more fuel efficient yet are more expensive, in combination with other cars. This tends to generate some 'stickiness' of vehicle demand for keicars in low density areas. For example, when the price of keicars increases, a consumer who prefers to own a combination of a keicar and a regular car may continue to own the same combination by accommodating this price increase by buying a cheaper regular car. This demand stickiness can reconcile both puzzles as it implies that low-income households in low-density areas may inelastically demand keicars. Hence, resolving the two puzzles essentially boils down to recovering the structural parameters of our demand model for vehicle ownership and utilization that vary by geographic space.

As with all analogous studies using household survey data [e.g., Bento *et al.* (2005), Bento *et al.* (2009), Goldberg (1998), Train (1986), West (2004)], the identification of the model parameters is challenging. We overcome this challenge by combining a set of control function approaches. First, to address the endogeneity of public transit density and rental prices in the vehicle ownership equation, we employ a version of the two stage residual inclusion (2SRI) method à la Terza *et al.* (2008) and Wooldridge (2015), using the 1980 railway networks as instruments. We further strengthen our identification by exploiting

³Japan is known for its highly efficient public transit network. Yet, public transit is only sparsely distributed in most non-urban areas. In Tokyo, for instance, only 31.7% of the working population drives to work (MLIT, 1998). This number is surprisingly small even compared to 65.7% in New York (FHWA, 2003). In Toyama, a moderately populated prefecture in Japan, the share of the working population who drives to work comes at 83.8% (MLIT, 1999). This number is close to what we observe in most U.S. counties.

 $^{^{4}}$ Keicar is an extremely small car segment with displacement levels of 660 cc or less. Keicars account for roughly 20% of the total vehicle sales in Japan.

the fact that the Japanese government implemented a variety of incentives for eco-friendly vehicles since 2001. We use information on purchase year/month to adjust the rental prices of *all* vehicles in each household's choice set. Second, to control for correlations between vehicle attributes and the error term in the vehicle utilization equation, we employ Dahl's control function approach. The use of Dahl's approach requires instruments that vary at the household level and affect vehicle purchase decisions, yet do not affect vehicle utilization decisions directly. In our case, this is trivially satisfied because the vehicle characteristics of *all* vehicles in the choice set enter the ownership choice while only those of the chosen car enter in the utilization equation. The exclusion restriction is further strengthened by the fact that we adjust the vehicle attributes of each consumer's choice set based on her purchase year/month. Lastly, in estimating the car utilization equation, we restrict the sample to those who purchased the cars only after 2012. This allows us to remove the spurious correlation that arise between our key variables and the unobservables due to the endogenous duration of car holding.

The estimated model not only resolves the two empirical puzzles mentioned above, but also produces a few new results that have not been documented (not extensively, at least) in the literature. First, our household-level analysis indeed signifies the important role of incomes, public transit, and portfolio considerations in explaining the geographic distribution of automobile demand. To demonstrate this point, we incrementally add to the demand model the interaction terms that capture (a) public transit density and (b) portfolio considerations, and compare each with the standard random-coefficient model that only accounts for non-specified preference heterogeneity. We see that our model improves the predictive power of the ownership shares of at least one car, hybrids, and keicars over geographic space by a large margin. Second, with all these factors taken together, we show that the estimated price elasticity of vehicle ownership is indeed smaller for keicars than for hybrid cars, overall as well as by public transit density. Thus, our household-level elasticity estimates help us explain the market-level estimates in Konishi and Zhao (2017). As we noted already, the key is that our model makes consumer's demand for automobiles 'sticky' for a certain subpopulation — consumers who prefer a certain combination of cars tend to stick to that combination in face of relative price changes. Third, Dahl's control function approach indeed works well in reducing bias due to sample selection. The polynomial sample-correction terms are jointly highly significant, and both income and price elasticity parameters have reasonable signs and magnitudes with the inclusion of these terms.

With the estimated demand system, we demonstrate the economic significance of spatial demand heterogeneity for efficiency and equity considerations. To accomplish this, we construct three counterfactual scenarios: (a) a carbon tax of \$50 per ton of carbon emissions,

(b) fuel-economy-based feebates (using the same social cost of carbon of \$50 per ton), and (c) a government-supplied hybrid-car-sharing platform. We also estimate policy impacts on consumer welfare using Herriges and Kling's (1999) approximation to the Small and Rosen's (1981) compensating variation formula. Our goal is to demonstrate the importance of spatial demand heterogeneity for economic evaluation, not to simulate realistic equilibrium responses to these policies. The latter is infeasible in our study because we lack data that would allow us to model spatially explicit supply-side responses. Hence, for each scenario, we simulate only the demand-side responses, assuming perfectly elastic supply (see **Subsection 7.1** for more discussion on this point).

Our simulation analysis deliver three important findings. First, carbon tax is far more effective in reducing on-road CO_2 emissions than feebates or ecocar sharing policies, yet the policy impacts different economic margins differently over different geographic regions. Second, there exists a large unexploited demand for *n*-household car-sharing, particularly in moderately dense areas, where the magnitude of the policy impact becomes roughly comparable to the carbon tax scenario. Third, carbon tax is estimated to induce a sizable welfare loss, more so in low-density areas than in high-density areas, whereas the ecocar-sharing policy is estimated to induce some welfare gain in a manner analogous to the consumer's gain from a new product innovation. These results indeed highlight the importance of accounting for spatial heterogeneity that arises from the intricate interaction of income, public transit, and portfolio preferences. Car owners in urban areas are generally richer, prefer having a smaller number of cars, and thus, have relatively inelastic demand for car ownership than those in non-urban areas. On the other hand, car owners in non-urban settings have inelastic demand for car utilization and have preferences for a mix of keicars with other cars, whose demand is estimated to be highly inelastic, precisely due to low public transit availability. Due to the mixed effects of these, each policy's impacts on car ownership and utilization tend to have a highly non-linear relationship to population density. Though not explored fully, our results also suggest the potential for a welfare-improving policy mix: a carbon tax in combination with a car-sharing platform in non-urban settings.

Our work complements three strands of literature: (a) empirical studies that investigate the relationship between urban structures and demand for vehicle transport, either using city-level observations [e.g., Levinson and Kumar (1997); Glasear and Kahn (2010)] or household-level observations [e.g., Beaudoin and Lawell (2018); Bento *et al.* (2005); Boarnet and Crane (2001); Train (1986); Gillingham (2014); Gillingham *et al.* (2015)]; (b) the economic incentives for efficiently controlling emissions from mobile sources [see Knittel (2012) and Anderson *et al.* (2011) for a comprehensive review on the topic]; and (c) empirical studies that estimate the discrete-continuous decision model on car ownership and utilization, with applications to the effect of gasoline tax [Bento *et al.* (2009), Train (1986), and West (2004)], to the effect of CAFE standards [Goldberg (1998)], and to the effect of feebates [D'Haultfoeuille *et al.* (2014)]. Our work is probably most closely related to Gillingham (2014) and Gillingham *et al.* (2015), which estimate the geographically explicit elasticities of driving with respect to gasoline prices in the U.S. context. To the best of our knowledge, however, none has explicitly examined the role of income, public transit, and portfolio considerations in understanding the spatial heterogeneity in demand for vehicle ownership and its implications for the design of a carbon-reduction policy or the environmental outcomes of car-sharing in a real empirical context.

2. A Statistical Overview of Vehicle Demand in Japan

We start by presenting a statistical overview of the relationship between access to public transit and vehicle-related household choice in Japan. For ease of visualization, we report household choice against a single composite index of public transit accessibility. We construct this index by an (unweighted) average of two district-level measures of railway transit network (incl. cable cars, surface rails, and subways). The first measure is the kilometers of railways per square kilometer and is intended to measure ease of access to destinations via railways. The second is the percentage of the habitable area within a district that has at least one train station within 15-min walking distance and is intended to measure ease of access to railways. We do not include the bus network in this index because in Japan, bus network is highly developed and even rural residents have access to a bus station within a walking distance. To confirm, **Figure 1-(A)** displays two scatter plots: the composite index and a similar index using bus network, both against the population density using districtlevel observations. The figure demonstrates that the composite index using railways has an increasing, but non-linearly relationship to population density while the index using buses has very little geographic variation and its inclusion would misleadingly overstate public transit accessibility. Hence, we use this index as the measure of public transit density.⁵

Figure 2 plots (A) the number of cars owned, (B) monthly vehicle kilometers traveled (VKT), (C) fuel economy ratings, (D) household income, (E) share of hybrid cars and "keicar", and (F) share of vehicle portfolio, all against public transit density.⁶ ⁷ In line with the

⁵We think that the frequency of service would give us a more reliable measure of public transit accessibility in case of bus service. Unfortunately, we do not have access to such data. However, we have a sense that the frequency of bus service would generate essentially the same geographic variation as our composite measure.

 $^{^{6}}$ We have detailed information on up to two most frequently used vehicles each household owns. Only about 5% of households own three or more vehicles. Hence, we use the sum over two cars for VKT and the (unweighted) average for fuel economy ratings, the hybrid vehicle shares, and the keicar vehicle shares. In figures are the averages of these over households.

⁷ "Keicars" are extremely small passenger vehicles with displacement of 660 cc or less. They are highly

theory and the findings around the world, vehicle ownership rate, VKT, and fuel-economy ratings decline sharply with public transit density. There is also a sign of self-selection: Only households with high enough demand for driving own cars, and thus, an average VKT tends to be high in the most dense area despite its public transit availability. Interestingly, however, demand for hybrid vehicles does not seem quite responsive to the density of public transit. Why don't households with limited access to public transit not own hybrid vehicles, despite their demand for driving seems twice as high as those in high density areas?⁸

A key to resolving this puzzle lies with the fact that there is a large gain in fuel economy by downsizing cars (Knittel, 2011). Households in low density areas tend to buy small-sized vehicles, that are both cheaper and more fuel efficient, for daily transportation. Thus, the share of keicars rises sharply with a decline in public transit density. Combined with the fact that average household income is lower in low density areas, the figure leads us to an observation that low-income households in low density areas are buying low-cost cars that offer high fuel-economy performance. This may still come as a surprise, though. Low density areas tend to have much wider roads and higher rates of traffic accidents. Thus, there should also be a gain for consumers from up-sizing cars. Resolving the puzzle, therefore, boils down to identifying the essential trade-off consumers make over fuel economy, size/safety, acceleration, and price of a vehicle, and how these trade-offs differ by income and public transit availability.

In considering this trade-off, our data also point to another important economic margin, which may become critical in identifying the demand parameters. Car owners rarely own two vehicles of the same type. Hence, households seem to make portfolio considerations in buying a vehicle rather than considering each vehicle in isolation. Since only 6% of households in our data own three or more cars, we consider the portfolio effect of two cars. In our data, of those who own two (or more) cars, only 36% own the same type of vehicles. The remaining 64% own a combination of either regular-keicar (24.5%), regular-minivan (16.3%) or keicar-minivan (23.2%). What seems critical for our empirical analysis is that such complementary portfolio holdings also vary sharply with the density of public transit. **Figure 2-(F)** demonstrates that the share of households who own *any* combination of different vehicle types increases

popular in Japan. Keicars account for roughly 30% of domestic car sales in Japan. The Japanese government offers a variety of tax incentives for these vehicles. As in **Figure 2**, we use the (unweighted) average over the two most frequently used cars for ownership share and car price.

⁸Some may argue that the share of hybrid cars is low in low density areas because the fuel economy of hybrid cars declines sharply in non-urban areas. This is not true, however. The Japanese hybrid cars do actually better in suburban conditions than in urban conditions. The primary disadvantage of hybrid cars is on highways — hybrids' fuel economy does decline sharply there. However, if this were the primary reason, then we should also expect low demand for keicars in low density areas — keicars do not do well either on highways either not only in terms of fuel economy but also in safety. But as shown below, demand for keicars rises quite sharply in place of hybrid cars. Hence, we must seek a different explanation.

as the public transit density declines, and the share of those who hold a keicar and another vehicle type (i.e., either regular or minivan) increases by even a larger margin. Whether this phenomenon can be explained purely by geographic or household-level variation or by explicitly accounting for portfolio preferences is another important empirical question, which we incorporate into our empirical investigation.

3. The Empirical Model

There is a large literature in transportation research that empirically examines consumer's vehicle ownership and utilization in the spirit of the continuous-discrete choice framework following Dubin and McFadden (1984). In that literature, consumer's choice is modeled as a two-stage decision process. In the first stage, the consumer chooses whether to own a car or not, and if she does, which type of car to own. In the second stage, the consumer chooses how much to drive over a given period of time. Our empirical framework follows this tradition, but extends it in several important ways.

Vehicle Ownership: We posit that consumers make a trade-off between money spent on buying a car versus the utility of owning and driving that car, as in the conventional literature [e.g., Bento *et al.* (2009), Berry *et al.* (1995, 1999), Goldberg (1998), West (2004)]. Thus, consumer *i*'s (indirect) utility from ownership of vehicle portfolio *j* in location *s* consists of two economic components, the expected utility from net income and the expected utility from vehicle ownership and utilization:

$$u_{ijs} = \rho \ln(y_i - r_{ij}) + v_{ijs}(X_i, Z_j, S_{is}) + \epsilon_{ijs}, \tag{1}$$

where y_i is *i*'s household income, r_{ij} is the annual rental price of vehicle ownership for car *j* for household *i*, and v_{ijs} is consumer *i*'s expected utility from owning alternative *j*, and ϵ_{ijs} is a pure stochastic error term distributed independently and identically across households, alternatives, and locations. v_{ijs} is a component that captures correlation across choices and heterogeneity across households due to household-specific attributes X_i , choice-specific attributes Z_j , or location-specific factors S_{is} , some of which are unobservable (stochastic).

Now, let us discuss our specification of the second ownership utility term v_{ijs} . Let j = 0 be an "outside option": i.e., not owning any vehicle. Naturally, consumers who choose this option would use public transportation for daily transport mode. Because this "index of desirability" v_{i0s} summarizes the maximal utility from sub-trip decisions conditional on owning no car (Domencich and McFadden, 1975), v_{i0s} (not u_{i0s}) should, in principle, depend on the

quality of public transportation (which is a component of S_{is}). Given the additive separability we assume in (1), we can re-define the second term as the *utility difference* relative to the no-car option, $v_{ijs} - v_{i0s}$, so the term now includes the value of public transportation.

We assume that $E[v_{ijs} - v_{i0s}|X_i, Z_j, S_{is}]$ has the following linear-in-parameter structure:

$$E[v_{ijs} - v_{i0s}|X_i, Z_j, S_{is}] = \delta'_{is}Z_j = (\delta_0 + \delta'_1X_i + \delta'_2S_{is})Z_j.$$

This specification conforms to a natural economic intuition that the marginal utility of a vehicle attribute varies by household as well as geographic characteristics. On the other hand, we can re-write this so it is expressed as the error-component formulation:⁹

$$\delta_{is}' Z_j = (\delta_0 + \delta_1' X_i) Z_j + \delta_2' S_{is} Z_j = \phi_{ij} (X_i, Z_j) + \xi_{ijs} (S_{is}, Z_j).$$

Thus, our econometric model is amenable to two interpretations that are mathematically equivalent. One is the random-coefficient interpretation: the marginal utility from a vehicle attribute depends on household-specific and location-specific factors (such as access to public transportation). Another is the error-component interpretation: the ownership value of a particular alternative j to a consumer depends on the consumer's underlying preferences for certain types of vehicles. Such consumer preferences depend naturally on householdspecific and location-specific factors because the ownership value, by definition, incorporates the value of sub-trips the consumer would make when she owns the car relative to the case of having no car (Domencich and McFadden, 1975). As discussed in Brownstone and Train (1998) and McFadden and Train (2000), this error-component structure can generate flexible substitution patterns (e.g., any type of nested logit as a special case), allowing us to alleviate the "independence from irrelevant alternatives (IIA)" property. Because we include a rich set of covariates in X_i and S_{is} , our model can flexibly capture sufficiently rich covariance structures of the error components, $E[\xi'_{ijs}\xi_{ijs}]$, underlying true substitution patterns. Furthermore, although the parameter on net income ρ is not allowed to vary by household, the income and price elasticities of demand still differ across households since $\partial u/\partial y = \rho/(y-r)$ and $\partial u/\partial r = -\rho/(y-r)$. These demand elasticities also depend on X_i, Z_j , and S_{is} , in general, because the consumer demand is a non-linear function of observables.

Portfolio Effect: Our discussion in **Section 2** signifies the importance of accounting for preferences for particular vehicle portfolios that may vary over geographic space. Hence, we augment the above model by allowing for the dependence of choices across multiple vehicle

⁹I am slightly abusing the term "error component" here. In the literature, the term "error component" refers to the stochastic component of the empirical model. In our specification, the second term ξ_{ijs} contains both stochastic and non-stochastic components.

holdings. Specifically, we follow Gentzkow (2007) and Wakamori (2015) and define consumer i's utility from owning a pair of cars j and k as follows:¹⁰

$$u_{i(j,k)s} = \rho \ln(y_i - r_{ij} - r_{ik}) + v_{ijs} + v_{iks} + \Gamma(j,k;X_i,S_{is}) + \epsilon_{i(j,k)s},$$
(2)

where $\Gamma(j, k; X_i, S_{is})$ is the portfolio-effect term, which captures the idea that households derive utility from owning a particular combination of vehicle types. For example, households with children may prefer owning a minivan for recreational use, yet may prefer owning a sedan or keicar for daily commuting use. As in Wakamori (2015), we consider three mutually exclusive sets of vehicle types: i.e., keicars \mathcal{K} , sedan/regular cars \mathcal{R} , and minivans \mathcal{M} . Then the portfolio effect is given by

$$\Gamma(j,k;X_i,S_{is}) = \kappa'_{(j,k)} x_{is},$$

where x_{is} denotes a vector of characteristics of household *i* in residence *s* (incl. constant) and $\kappa_{(j,k)}$ is the combination-specific parameter for a pair (j, k).

There are several advantages of modeling the portfolio considerations this way. First, as discussed in Wakamori (2015), the approach does not assume products are either complements, substitutes, or independent, and instead, allow the estimates of parameters to flexibly capture complementarity patterns observed in the data. Second, we can estimate this model using conventional conditional logit routines available in most statistical packages. Lastly, this formulation exploits an important property of mixed logit: an analog to nested logit of any complexity can be obtained by adding interaction terms with a set of dummies representing the nests (Brownstone and Train, 1998). In eq. (2), we are just adding a set of dummies, each representing a particular portfolio, and then, interacting each of these dummies with household-level or geographic-level observables. The former essentially works the same as having a nest for each vehicle portfolio while the latter works as allowing the correlation across choices within the nest to depend on observables.

Vehicle Utilization: Following the convention [e.g., Bento *et al.* (2005), Bento *et al.* (2009), Goldberg (1998), West (2004)], we assume that monthly driving distance m (in log) of consumer i who lives in area s and who owns vehicle j is:

$$\ln m_{ijs} = \alpha_{is} \ln(y_i - r_{ij}) + \beta_s \ln p_{ij} + \lambda' W_{ijs} + \eta_{ijs}, \qquad (3)$$

where y_i and r_{ij} are as defined above, p_{ij} is the operating cost of utilization per unit of driving

¹⁰In this study, we restrict consumer's choices to two vehicles per household since we have detailed information only on two most frequently used cars. As discussed in **Section 2**, only 6% of the households in our data hold three or more cars.

distance for car j for household i, W_{ijs} is a vector of household, vehicle, and geographic characteristics, and η_{ijs} is the error term. The primary parameters of interest are the income and price elasticities, α_s and β_s , of vehicle utilization. We allow these parameters to depend on the geographic characteristics S_{is} :

$$\alpha_s = \alpha + \gamma'_{\alpha} S_{is}; \quad \beta_s = \beta + \gamma'_{\beta} S_{is}$$

In contrast to eq. (1), eq. (3) is the (reduced-form) equation, and thus, without explicitly allowing for interaction terms, the income/price elasticity will be constant across areas.

It is known that OLS regression of eq. (3) would generally give us biased estimates of parameters due to sample selection because we observe each consumer's driving behavior only for the car model chosen, but not for car models that had not been chosen. To see this in our empirical setup, note that we can write $v_{ijs} = E[v(m_{ijs})|X_i, Z_j, S_{is}] + e_{ijs}$, and hence, the error term ϵ_{ijs} in eq. (1) is confounded with another error term e_{ijs} , forming the joint error term $\mu_{ijs} \equiv e_{ijs} + \epsilon_{ijs}$. Consequently,

$$E[\eta_{ijs}|X_i, Z_j, S_{is}, j \text{ is chosen}]$$

= $E[\eta_{ijs}|V_{ijs} + \mu_{ijs} \ge V_{iks} + \mu_{iks} \text{ for all } k] \neq 0$

where V_{ijs} is the observable part of the indirect utility, and the last inequality follows because η_{ijs} contains some of the information in e_{ijs} , the unobserved part of utility from driving car model j. Simply put, consumers would enjoy driving cars of their favorites and not so much for others.

To address this selection problem, previous studies either assumed a joint distribution of errors (μ_{ijs}, η_{ijs}) or used a selectivity correction à la Dubin and McFadden (1984). The former is known to place severe restrictions on the selection process, while the latter is known to become imprecise or infeasible when there are many alternatives in the first stage decision. We instead use Dahl (2002)'s control function approach to correct for this selection bias in the case of many alternatives. Specifically, Dahl showed that, in case of high-dimensional alternatives, eq. (3) can be consistently estimated using estimates of individual purchase probabilities:

$$\ln m_{ijs} = \alpha_s \ln(y_i - r_{ij}) + \beta_s \ln p_{ij} + \lambda' W_{ijs} + \sum_{j=1}^J M_{ij} \times T_{ij}(P_{i0}, P_{i1}, \dots, P_{iJ}) + \upsilon_{ij}, \quad (4)$$

where $T_{ij}(\cdot)$ is some unknown function of purchase probabilities P_{i1}, \ldots, P_{iJ} and M_{ij} is its parameters. Dahl suggests that, in practice, we may include only a few probabilities such as

the probabilities of the first-best choice, the second-best choice, and the outside option. We follow this advice in our estimation.

We can incorporate the portfolio effect into the sample correction terms in the vehicle mileage equation (4) by slightly modifying our notation. Let J_1 and J_2 be the sets of products for her first and second cars, respectively. Let us augment J_2 by including 'zero', an option to own no second car. Adjoining these two sets and an outside option to own no car, we create the joint choice set J, which contains $1 + \#J_1 \times \#J_2$ alternatives. The consumer chooses an alternative j from this adjoined set. That is, one may choose to own no car (0,0), choose to own one car $(j_1,0)$, or choose to own two cars (j_1, j_2) . With a slight abuse of the notation, (j, k) in place of j, the model described by (1) and (4) is essentially intact.

There is one subtle, yet important, issue in estimating the VKT regression (4) — we observe VKT for *each* of the vehicles owned, and hence, (4) must be estimated separately for each vehicle, accounting for that vehicle's attributes. Here, the difficulty is that the households who own multiple cars are likely to decide on how often to use one vehicle jointly with other vehicles. Consequently, the utilization levels are likely to be correlated across car holdings. The literature to date seems silent as to how to address this issue. In her seminal work, Goldberg (1998) estimates the VKT regression using observations on newly purchased cars only, ignoring this correlation in vehicle utilization. Bento *et al.* (2005) instead use the VKT per vehicle, averaged over vehicles owned, as a dependent variable, excluding vehicle-specific regressors from the list of independent variables.

In the context of the present paper, the policy impact on the second car's VKT is quite important. Hence, we address this issue as follows. We estimate the VKT regression, pooling all VKT observations on the two most frequently used cars, with a dummy indicating a second car. This ensures that the same sample correction terms enter the VKT regression for the two cars owned by the same household, yet accounting for the fact that one of the observation is on the second car. This allows multiple-car owners' vehicle utilization decisions to be correlated across their vehicle holdings, either through observable householdlevel characteristics or through (unobservable) selectivity terms. Because our model of car ownership accounts for portfolio effects, the selectivity correction terms in (4) control for the unobservable correlations that are specific to the same household who decide to own a particular combination of cars.¹¹

¹¹One could, instead, estimate the seemingly unrelated regression or the second-car's VKT equation independently. Both approaches resulted in parameter estimates that are hard to interpret. For example, the estimated elasticity on net income was negative. We would think that this occurs precisely because of the substitution in vehicle utilization between the two cars. The households with high incomes primarily drive the first car for daily use, keeping the second car only for luxurious use. The households with low incomes, on the other hand, are likely to own the second car for primary use, and hence, they drive the second car more. The estimates may be simply capturing this correlation.

4. Data

Our study relies on a large cross-sectional sample of households from a nationwide internet survey conducted in November 2016 in Japan. In designing the survey, we aimed for two goals. The first is to obtain a sufficiently large sample, with variations in household-level characteristics, for each population density decile. This is essential for our study because we need comparable households to separately identify the demand parameters that vary by income and public transit: i.e., those with different levels of income, yet with the same level of access to public transit as well as those with the same level of income, yet with different levels of access to public transit. The second is to collect sufficiently detailed information on each household's vehicle ownership and utilization that is comparable to the U.S. Consumer Expenditure Survey (CES). In particular, we aim to collect information such as the number of vehicles, the vehicle type (fuel economy, engine/fuel type, horsepower, make, size, weight, vintage), the year/month of purchase, and the vehicle kilometers traveled since the purchase. Such detailed information on vehicle ownership and utilization is not available in national consumer surveys in Japan.

The survey was administered under the contract with Nikkei Research Inc. to the pool of registered internet monitors. The survey resulted in a sample of 105,000 usable respondents with complete responses. As with other internet-based surveys, we did not have direct control over the sampling process. Our usable sample, however, covers a sufficiently large number of households in every prefecture, with sufficient variation in key socioeconomic variables such as age and income. In the **Online Appendix (A)**, we compare our sample distribution against the population distribution by prefecture. The geographic distribution of our survey respondents by prefecture is sufficiently close to the population distribution, though populated prefectures (e.g., Tokyo and Kanagawa) are over-represented while less populated prefectures in Kyushu region) are under-represented. As expected, average household incomes in our sample are slightly higher than in the population for most prefectures, although we do not see significant differences in average household sizes. Our results may be somewhat biased toward households with relatively higher incomes.¹²

We supplement the survey with the data from various sources. First, we use the GIS datasets on city boundaries, bus stops, train stations, train networks, hospitals, road length, and public parks from the National Land Numerical Information Download Service, made available online by the Ministry of Land, Infrastructure, and Transportation (MLIT). We

 $^{^{12}}$ There is a large literature in environmental economics, examining the extent of bias in demand estimation that may arise due to the internet-based survey. The results are mixed. Comparing the internet survey versus other modes of survey, some (Lindhjem and Navrud, 2011 and Nielsen, 2011) report no or small bias while others (Boyle *et al.*, 2016) report a non-negligible bias.

use the coordinates of train stations and the line data on train networks to construct the composite measure of public transit density at the 'city-district' level (see Section 2). Our definition of 'city-district' follows that of the Ministry of Internal Affairs and Communications (MIAC). As of 2018, there are 1,724 city districts in Japan. Second, we also use the car catalog data from the carsensor.net, one of the largest online car retailers in Japan. The survey respondents are asked to provide detailed information on each of the cars they own (up to their second car): i.e., model year/month, purchase year/month, make, model name, displacement level, curb weight, and mileage. We use these to match their cars with those listed in the carsensor catalog to obtain other vehicle characteristics such as fuel economy ratings, horsepower, size, and transmission. Third, we use the district-level population estimates provided by the National Institute of Population and Social Security Research (IPSS). Lastly, we also obtain a measure of prefecture-level road congestion from MLIT, regional consumer price and gasoline price series from the Ministry of Economy, Trade and Industry (METI), historical discount rates from the Bank of Japan, and district-level garage certification regulations from Keicar Information Center. Detailed descriptions on how we define our choice set (for vehicle ownership) and key variables used in the manuscript are available in the **Online Appendix** (B).

Table 1 reports the means and standard deviations of key variables by population density. The table confirms substantial variations both within and across population density quintiles, which we exploit in our estimation. First, all measures of public transit sharply decrease as population density declines. Not only that, we have substantial variation in these measures within each density quintile, and interestingly, more so in low-density quintiles: coefficients of variation are 0.57, 0.45, 0.45, 0.28, and 0.21 for the lowest, 4th, 3rd, 2nd and the highest population densities. This is in sharp contrast to household characteristics. Average household incomes decline as population density declines, yet the coefficients of variation stay roughly the same across all quintiles. The same is true with household size. This 'within' variation in public transit measures helps us identify the effects of public transit on vehicle ownership/utilization. Second, as we have seen, the rate of car ownership rises quickly as population density declines, possibly in response to declines in public transit availability. Interestingly, however, the coefficient of variation for car ownership declines as population density declines. Instead, the coefficient of variation for the number of cars owned rises, from 0.33 in the highest density quintile to 0.52 in the lowest density quintile. This point is also closely related to our next observation. Third, we observe a smaller variation in vehicle utilization than vehicle ownership: after taking logs, the coefficients of variation for monthly VKT range from 0.17 to 0.19 for the first car (= most frequently used car), and from 0.18 to 0.24 for the second most used car. Combined, these two observations are suggestive of the tendency that households absorb the impact of public transit (un)availability by adjusting the number of cars owned rather than by adjusting the vehicle utilization. We take this as suggesting that it is indeed important to account for correlation between ownership decisions and utilization decisions. Lastly, the sample characteristics of the first car seem to differ substantially from those of the second car. The second cars are cheaper, more fuel-efficient, and smaller on average in virtually all density quintiles (while coefficients of variation are similar between the first cars and the second cars). Interestingly, for their first cars, households in low density areas are more likely to own hybrid cars than in high density areas. Yet, the opposite is true with their second cars. These points seem to re-confirm the existence of the portfolio effect discussed in **Section 2**.

5. Estimation and Identification Strategy

The discrete-continuous choice model we develop in **Section 3** is estimated in two steps. In the first step, we estimate the discrete choice model, assuming the form of indirect utility as in eq. (2) and the Type-I extreme value distribution for ϵ . For this step, estimation is done by Stata's alternative-specific conditional logit routine. In the second step, we estimate the VKT regression in eq. (4), pooling all VKT observations for all cars owned by households in the sample. In this step, we use polynomials of predicted probabilities from the first step as selection control terms as in Dahl (2002). We experiment with a polynomial of up to third degree, using the probabilities of the chosen, the no-car, the highest-likelihood, the second highest-likelihood, and the lowest-likelihood options. Based on the sign/significance of key variables (i.e., net income and cost per kilometer of driving) as well as their robustness to varying levels of controls, we end up using the second-degree polynomials of the highest, the second-highest, and the lowest probabilities. Furthermore, with this approach, the conventional covariance estimator is biased (Dahl, 2002). Hence, we use bootstrapped standard errors, with 500 draws, for inference.

Because we use one-shot household survey for both steps, the identification of the parameters relies on cross-sectional variation at both the household and the district levels in economic/geographic variables. Though this poses a challenge in identification, this is typical of studies that estimate the discrete-continuous choice model of car-holding decisions using survey data (see Goldberg, 1998, Bento *et al.*, 2005, and Bento *et al.*, 2009). In the literature, four identification challenges are discussed: (1) endogeneity of measures of public transit and (2) endogeneity of rental price of car ownership in the first-stage choice of car ownership; (3) endogeneity of operating cost of car utilization (due to sample selection) in the second-stage choice of car utilization. In addition to these, there is an issue with the

endogenous duration of car ownership. Below, we discuss each of these and our strategies to address them in order.

Public transit: In the literature, it is often assumed that public transit is predetermined prior to their car-holding decisions. Indeed, these measures are known to work better than other forms of urban structures (see Bento et al., 2005). However, there is still a concern that households' intrinsic preferences for car holdings may be correlated with measures of public transit — households may make residential choice in conjunction with choice of car holding. To address this concern, we exploit the idea that the public transit network in the past is generally a good instrument (e.g., Duranton and Turner, 2011). There are, however, two limitations with the conventional IV approach in this context. First, the correlation between the public transit measure and the unobservable errors makes all of its interaction terms endogenous. Hence, the conventional IV strategy would require a large number of instruments. Second, it is known that a 'plug-in' 2SLS method produce inconsistent estimates in nonlinear models such as this. Terza et al. (2008) and Wooldridge (2015) discuss how an alternative two-stage residual inclusion (2SRI) method can overcome these limitations. In our estimation, we further exploit the parsimonity of the 2SRI approach using the 1980 railway networks as instrument (see **Figure 1-(B)**) to addressing virtually all endogeneity concerns that arise through consumer's endogenous residential decision. Our **Online Appendix (C)** explains how this is accomplished.

Rental price of car ownership: Studies on the automobile demand estimation are often concerned with the endogeneity of car prices. There may be product attributes consummers see but researchers do not, such as brand images, styles, and non-price incentives. Since they are demand-shifters, they may as well be correlated with car prices. This concern is less serious in our study because ours is based on household data and car prices are mostly determined at the market level. In addition, we include make, car-type, fuel-type, and usedcar dummies to control for unobservable product characteristics. This identification strategy is analogous to Goldberg (1998) and Bento et al. (2009). However, individual households also negotiate prices at the dealer level. Hence, there may be some measurement error in our price variable that may be correlated with vehicle attributes at the local level. To take care of the concern, we use *time-varying* car-related tax incentives as exogenous price shifters. Specifically, we use information on purchase year/month to adjust the rental prices of all vehicles in each household's choice set. The Japanese government implemented a variety of incentives for eco-friendly vehicles since 2001 [see Konishi and Zhao (2017) for more detail]. This not only gives us exogenous price variation, but also implicitly restrict each household's choice set.

Sample selection: For the second-stage choice of car utilization, we essentially have two

identification issues. Both issues are closely related to the endogeneity of p, i.e., the operating cost of vehicle utilization. The first issue concerns the sample selection we discussed in **Section 3**. Because the fuel cost per kilometer is a function of the fuel economy rating of the car chosen, this variable is clearly correlated with unobservable demand for driving distance. To control for this, we use Dahl's control function approach. This approach requires the exclusion restriction: i.e., we need instruments that vary at the household level and affect vehicle purchase decisions, yet do not affect vehicle utilization decisions directly. In our case, this is trivially satisfied. Vehicle utilization in eq. (4) depends only on the attributes of the car that is actually owned, but not those of alternatives, while vehicle ownership choice in eq. (2) depends not only on the attributes of the chosen car but also on those of the other alternatives. Previous studies essentially use the same argument in implementing Dubin-McFadden type correction. The identifying condition is further strengthened by the fact that we use the prices of alternative vehicles actually observed at the time of purchase instead of using contemporaneous values observed today, as discussed above.

Duration of car ownership: The second identification issue for the vehicle utilization regression concerns the duration of car ownership. The problem here is that those who own cars longer tend to be those with low incomes and own cars with low fuel economy ratings (both because cars sold in the past tend to be fuel-inefficient and because fuel-efficient cars tend to be expensive), while at the same time, these households continue to hold cars despite their low incomes precisely because they have high demand for driving. Hence, households who hold the same car for a long time tend to be those with lower incomes and higher costs of vehicle utilization. This results in spurious correlations that bias the parameter estimates in the opposing direction — i.e., negative correlation between net income and VKT and positive correlation between the cost of driving and VKT. Our first-stage model of car ownership allows us to account for the economic margins that affect "whether or not", "how many", and "what type of cars" to own, but not "how long" to own. Naturally, the selection correction terms cannot control for the endogeneity that arises from the duration of ownership. To address this concern, we restrict the observations to those on cars purchased after 2012.

6. Estimation Results

6.1. Vehicle Ownership

We first report on the first-stage discrete choice model of car ownership in **Table 2**. Three sets of results are reported in the table: The models without portfolio effects, with portfolio effects, and with portfolio effects and the 2SRI controls. For each model, the first row exhibits the estimates of the mean parameters for our key variables while the second and the third columns present their interaction effects with transit availability and household size. These interactions allow us to account for cross-region as well as within-region heterogeneity in demand for car ownership. The estimated model also includes make dummies (Toyota, Honda), a used-car dummy, fuel-type dummies (hybrid, diesel), vehicle-type dummies (keicar, minivan), a garage certificate requirement dummy as well as their interactions with metropolitan dummies to tease out the effects of unobserables.

Virtually all parameters are statistically highly significant, and their signs are consistent with economic theory as well as previous studies that estimated similar models. First, the parameter on the logged net income is significantly positive, which implies that consumers with higher incomes are more likely to own cars and that consumers prefer cheaper cars, holding all else constant. Second, the mean parameter on the fuel cost per kilometer (Japanese yen per km, YPK) is significantly negative, which implies that consumers on average value fuel economy. However, its interaction terms suggest that consumers with access to transit density or with large family tend to care less about fuel economy (even after controlling for vehicle size and vehicle types). Third, consumers on average prefer high acceleration. Interestingly, consumers with large family size tend to value acceleration much less. Lastly, though not reported, the keicar dummy is significantly positive, whereas diesel and hybrid car dummies are significantly negative. The Japanese consumers thus prefer keicars over regular gasoline cars, yet prefer gasoline cars over diesel or hybrid cars. This occurs because in Japan, diesel cars are not popular as they are often perceived as unsafe (due to its low ignition temperature) or unclean (due to its high sulfur content before desulfurization process).

There is one anomaly that may seem at odds with previous studies in the U.S. — the mean parameter on car size is significantly negative, implying that in Japan, consumers on average value smaller cars. Note, however, that we obtain these estimates *after* controlling for the vehicle types (i.e., keicars, minivans etc) and portfolio effects, which also vary by household size. Hence, the negative sign on car size should be capturing the preferences for the compactness of vehicles within, but not across, vehicle class. Our interpretation therefore is that because the roads and parking spaces are narrow virtually everywhere in Japan (even in rural areas compared to roads in U.S.), consumers on average prefer smaller cars, given their preferred vehicle class. This is consistent with the findings in Konishi and Zhao (2017).

Table 2 also signifies the importance of accounting for portfolio considerations — both mean and interaction parameters on many of the portfolio terms are statistically significant. The mean parameters on *keicar* combinations (except on *keicar-keicar*) are significantly positive, suggesting that consumers, on average, value the *keicar* combinations more than

the regular-regular combination, which we take as the base combination. However, their interaction terms with public transit is significantly negative and economically large. It implies that (combined with the household size interaction terms) average households in non-urban areas value the *keicar* combinations *more* than the base combination. Thus, these estimates generate the stickiness of consumer demand to *keicars* in non-urban areas, either held as a single car or as a combination with other types, confirming the empirical patterns we observe in Section 2. Essentially the same goes for minivan combinations (regular-minivan, minivan-minivan). The mean parameters on minivan combinations are positive and statistically significant. Yet, their interaction terms with public transit and household size are statistically significant and have the same signs and magnitudes as keicar combinations. The results are consistent with the idea that consumers adjust their vehicle portfolios according to the households' needs. It may seem counter-intuitive, however, that the value of any combination (relative to the base) decreases with household size. But this is only after conditioning on the value of vehicle size, which is estimated to increase with household size. Thus, we interpret the result as suggesting that consumers with a larger household size do prefer owning a larger car, but do not necessarily prefer a family-car portfolio (e.g., *minivan-minivan*) over a regular-car portfolio.

To gauge the importance of accounting for preference heterogeneity over geographic space, we also compare the predicted shares from alternative empirical models against the observed shares by transit density. The first is our full model (the model with spatial interactions and portfolio effects) and the second is the conventional random-coefficient (RC) logit (the model that does not account for spatial heterogeneity or portfolio effects). **Figure 3** reports the results of this exercise for three types of ownership: Ownership of any car (panel A), hybrids (panel B), and keicars (panel C). Hence, the figure evaluates the predictive performance on two economic margins: whether or not to own a car and which type of car to own. Note that in these figures, we report on the *unconditional* ownership shares — i.e., not *conditioned* on having a car.

As shown in the figure, the naive RC logit fails to predict all ownership shares by a large margin. We also see a large swing in the prediction errors — i.e., it tends to understate all types of ownership shares in low-density areas whereas overstating them in high-density areas. In contrast, our full model predicts the car ownership share quite precisely for all transit density levels. Importantly, the panel B and C demonstrate that our model does a far better job of predicting the ownership shares of hybrids and keicars, allowing us to explain the empirical patterns discussed in **Section 2**. The accuracy of prediction on these two margins is quite important in simulating the counterfactuals policies in **Section 7**.

6.2. Vehicle Utilization

Next, we turn to the vehicle utilization regression. Our focus is on demand elasticities with respect to (net) income and operating cost, and on the influence of public transit density on these elasticities. The first four columns of **Table 3** report the results without the congestion variable interacted with the key variables: income and operating cost (in log). The next four columns include these interaction terms. We estimate each of these regressions with varying levels of controls: metropolitan dummies and selection correction terms. Though the parameter estimates are not reported, all regressions control for other vehicle characteristics (fuel type and car type dummies), demographic characteristics (age, household size, marital status, number of cars owned, work status, distance to work, years of education) and urban structures (district-level population density, access to hospital, and access to public parks). As discussed in **Section 5**, we only report the results with the second-degree polynomials of the highest, the second highest, and the lowest estimated choice probabilities for selectivity correction.

The estimate of the mean income elasticity is always positive, but become larger and more statistically significant when we include the congestion interaction term. This makes intuitive sense. When a household's income increases, the household would increase her time to allocate for leisure, but how much she would increase time to spend on driving depends on how congested roads are. She would drive more if roads are less congested. Furthermore, the income elasticity is smaller for consumers living in high density areas. This makes sense since recreational value of driving would be larger for consumers who have limited access to in high public transit. Note that we can focus on leisure-related arguments as we already control for distance to work.

The estimate of the mean price elasticity is negative and statistically significant across all specifications. This is consistent with economy theory, but suggests the success of our control strategy — in studies that use cross-sectional household-level data, this elasticity is often estimated with bias toward zero or even positive [e.g., Goldberg (1998)]. The interaction with public transit density is positive and marginally significant. This suggests that the demand for driving is less price-elastic in high density areas. This may seem somewhat counter-intuitive at first. When the price of gasoline increases, for example, consumers who have access to public transit can use public transit instead of driving, and therefore, we would expect the demand for driving to be more price elastic in areas with high public transit use cars primarily for a recreational purpose, and the fuel cost accounts for a relatively small portion of the recreational expenditures. On the other hand, consumers with limited access to public transit use cars for daily use, and the fuel cost accounts for a larger share of consumer

expenditures. Hence, this makes the demand for driving more price-elastic for consumers in areas with low public transit density. The inclusion of the congestion interaction term makes the interaction term with transit density become less statistically significant and smaller in magnitude. Our reasoning is that congestion and transit density tend to co-move, and thus, are picking up similar effects.

With all specifications, the selectivity correction terms are jointly highly significant. This is despite the fact that we already control for a number of observables. Moreover, inclusion of these correction terms generally improves the statistical significance of our key parameters. This implies that there is selection on unobservables, and thus, omitting selection correction terms is likely to bias our estimates. When translating these parameter estimates into elasticity estimates by public transit quintile, this gain in consistency does seem to matter. Hence, we are in general in favor of models with selection correction. On one hand, the estimate of the mean income elasticity parameter is positive and gets larger in magnitude after selection correction. This means that unobservable demand factors for vehicle utilization, which the selectivity terms are meant to control for, are negatively correlated with household (net) incomes. On the other hand, the estimate of the mean price elasticity parameter does not seem to change much, both in magnitude and in statistical significance.

6.3. Elasticity Estimates

Table 4 reports the estimates of various elasticities of car ownership and utilization by transit density quintile. We estimate the price elasticity ϕ_{dt} of car ownership with respect to rental price r for fuel type t (t = hybrid, diesel, keicar) for transit density quintile d as follows:

$$\phi_{dt} = \frac{\partial s_{dt}}{\partial r_{dt}} \cdot \frac{r_t}{s_{dt}}$$

where s_{dt} and $\frac{\partial s_{dt}}{\partial r_t}$ are calculated as

$$s_{dt} = \frac{1}{N_d} \sum_{i \in I_d} \hat{P}_{it} \text{ and } \frac{\partial s_{dt}}{\partial r_{dt}} = -\frac{1}{N_d} \sum_{i \in I_d} \frac{\hat{\rho}}{y_i - r_{it}} \hat{P}_{it} [1 - \hat{P}_{it}],$$

where \hat{P}_{it} denotes the estimate of household *i*'s probability of holding car models of fuel type t, $\hat{\rho}$ is the estimate of the parameter on net income, $y_i - r_i$ is household *i*'s observed net income, N_d is the number of sample households in d, and I_d is the set of households in d. For the price elasticity of VKT, we simply evaluate the parameter estimates $\hat{\alpha}_d$ and $\hat{\beta}_d$ for $\ln(y - r)$ and $\ln(YPK)$ at the means of congestion and transit density measures for each transit density quintile d, using the estimates from our preferred specification (8) in Table 3. The price elasticities of car ownership range from -0.223 to -1.009. The numbers are in line with, but slightly lower than, those reported in analogous studies (Bento *et al.*, 2009; Konishi and Zhao, 2017).¹³ As expected, the car ownership elasticities get smaller (more inelastic) with transit density — consumers in low density areas are more price elastic than those in high density areas. Moreover, hybrid cars face the most elastic demand. In contrast, keicars face the most price inelastic demand. We believe this is the key to resolving the opening puzzle. Consumers in low density areas have lower incomes, and hence, are more price elastic. Yet, even these consumers have substantially more price elastic demand for hybrid cars than keicars. Consequently, consumers in low density areas tend to buy keicars for cheaper prices despite their high level of vehicle utilization. This also offers support for the earlier findings of Konishi and Zhao (2017) who also find that keicars face more price inelastic demand, despite that keicars typically serve low-income groups whose demand tend to be more elastic *ceteris paribus*.

Vehicle utilization is not quite elastic with respect to (net) income. Interestingly, the income elasticities decline with public transit density, and turn practically zero (i.e., statistically highly insignificant) in the forth to first transit density quintiles. We believe this reflects the competing effects of income. As consumer's income rises, the consumer increases the time to spend on leisure, which tends to increase vehicle utilization, but at the same time, it also increases the opportunity cost of time. With the high congestion level in urban areas where public transit availability is also high, the consumer faced with the high opportunity cost of time tends to opt for public transit rather than drive their cars. In Japan, roads in even moderately populated areas are still congested, compared to the U.S. Hence, it is quite reasonable to observe that the income elasticities in such areas tend toward zero (or even negative).

The price elasticities of VKT are negative and statistically significant in all transit density quintiles, but slightly lower in magnitude than those reported in Bento *et al.* (2009). Because our car utilization regression focuses on a subsample of households who have bought cars since 2012, much of the variation in YPK comes mainly from cross-sectional variation in fuel economy ratings rather than that of gasoline price. Hence, our price elasticity estimate is essentially capturing the rebound effect: i.e., buying a more fuel efficient car makes a consumer drive more. Our results are consistent with the argument that the rebound effect is often overstated (Gillingham *et al.*, 2013). Our estimates suggest that the rebound effect is even smaller in Japan than in the U.S. However, the use of cross-section data is known

 $^{^{13}}$ Note that we take the 'holding' approach as opposed to the 'transactions' approach in our demand estimation [see Goldberg (1998) or Bento et al. (2009) on the difference between the two]. With this approach, we implicitly get at the long-run demand elasticities, which tend to be smaller than the short-run elasticities.

to bias the price elasticity estimate toward zero due to the selection bias [for example, Goldberg (1998) obtain a positive and statistically insignificant price elasticity estimate]. Hence, there is a possibility that our elasticity estimates may still be biased toward zero even after controlling for the selectivity terms.

7. Counterfactual Analysis

7.1. Counterfactual Scenarios

We now use the estimated demand system to simulate the spatially-explicit distributional impacts of counterfactual carbon-reduction policies over geographic space. There is already a large literature that empirically investigates the distributional consequences of automobilerelated economic policies [e.g., West (2004), Bento *et al.* (2009)]. Their focus is, however, on the variation of policy impacts by income. In contrast, what we wish to demonstrate here is that public transit availability and innate preferences for vehicle portfolios may interact with income distribution over geographic space in such a way that generates important implications for efficiency and equity of alternative climate mitigation policies. To that end, we consider three counterfactual policies: (a) carbon tax of \$50 per ton, (b) feebates based on fuel-economy ratings of cars, and (c) a public provision of an ecocar-sharing platform.

Our outcome variables of interest are vehicle ownership (number and types of cars owned), vehicle utilization (monthly vehicle kilometers traveled), carbon emissions (monthly CO_2 emissions from vehicle utilization), and consumer welfare (to be defined below). Our ability to simulate economic outcomes is limited to those that make use of demand-side parameters only. We have neither data nor policy relevant variations to estimate the supply-side parameters with respect to the public transit or the car-sharing platform. Therefore, in the analysis below, we assume perfectly elastic supply-side responses for all policy scenarios. Consequently, we do not attempt to simulate car prices, gasoline prices, and any other economic outcomes that would require supply-side parameters. These are simply treated as 'fixed' in our simulation analysis. We also assume no revenue recycling (or revenue accounting) because it would only obscure the essence of the main analysis. Consequently, our simulated outcomes should not be taken as realistic equilibrium responses to these counterfactuals. Our goal here, instead, is to conduct a quantitative evaluation of how income, public transit density, and portfolio preferences may interact to generate heterogeneous demand-side responses over geographic space, signifying the importance of accounting for such spatial heterogeneity for designing efficient and equitable climate mitigation policies.

Counterfactual 1. Carbon Tax

In this policy scenario, we assume that a carbon tax is implemented on top of the existing tax on gasoline. This is indeed consistent with other countries' experiences [see, for example, Andersson (2019) for the case of Sweden]. We use \$50 per ton of CO_2 as a benchmark estimate of social cost of carbon emissions (SCC) per Revesz et al. (2017).¹⁴ In theory, this carbon tax on gasoline can fully restore economic efficiency by correcting for the negative externality associated with carbon emissions from vehicle transportation. Hence, this policy scenario serves as an important benchmark to contrast other policy counterfactuals. In our empirical model, the carbon tax raises operating costs p_i of all cars, yet it, in general, raises the operating costs of fuel-inefficient cars (non-keicars) more than those of fuel-efficient cars (keicars). Hence, this policy is expected to shift consumer demand from less fuel-efficient cars to more fuel-efficient cars while decreasing overall vehicle ownership and utilization. Our focus here is the distribution of such responses over geographic space and its implications for consumer welfare vis-à-vis other policy scenarios. There is an important question as to how the rental prices r_i respond to the changes in p_i . In our simulation, rental prices stay unchanged since we assume perfectly elastic supply, which would absorb all the price impact of the demand shocks.¹⁵

Counterfactual 2. Feebates (Eco-car Tax/Subsidy Incentives)

Next, we consider eco-car incentives on car holdings. Tax and subsidy incentives are a common policy apparatus for inducing purchase of eco-friendly cars around the world. Such incentive schemes are more generally termed as 'feebates' since they impose fees (or taxes) on less fuel-efficient vehicles and rebates (or subsidies) on more fuel-efficient vehicles (See Anderson *et al.*, 2011). Institutional details on these incentive schemes vary by country. A number of empirical studies have credibly quantified the economic impacts of such incentive policies in a variety of contexts: the U.S. (Beresteanu and Li, 2011), France, Germany, and Sweden (Klier and Linn, 2015), and Japan (Konishi and Zhao, 2017). Our purpose here is not to repeat these studies; rather, to understand how consumer's response may vary over geographic space due to the intricate interaction of income, public transit access, and portfolio considerations. Hence, we do not take any particular country's incentive scheme,

¹⁴Recent advances in the SCC literature suggests a much higher estimate of SCC [see, for example, a report by the Energy Policy Institute at the University of Chicago (EPIC, 2021)]. We experimented with various SCC values and confirmed that much of our qualitative discussion stays the same, although the magnitudes of policy impacts change.

¹⁵Alternatively, we may either assume that rental prices would fall exactly by the net present value of a corresponding increase in annual fuel costs [per Busse *et al.* (2013) or Allcott and Wozny (2014)], or assume the Bertrand-type price competition and solve for new equilibrium prices, in which case the changes in rental prices depend on each automaker's market power. In either case, the resulting automobile ownership would be higher. Hence, our result may be taken as the upper bound estimate of the policy impact.

and instead, consider a general feebates scheme: tax and subsidy in proportion to its carbon footprint. Specifically, the tax t_j (or the subsidy) is charged to the holding of car j according to the formula:

$$t_j = \tau \times E[v] \times (EPK_j - x),$$

where τ is the social cost of carbon emissions, E[v] is the expected annual VKT, EPK_j is the carbon emissions per kilometer of driving distance, which is estimated as the emissions per liter of gasoline divided by the fuel economy (km/L), and x is the base emissions rate beyond which consumers pay a tax and below which consumers receive a subsidy. As with the carbon tax, we use the SCC value of \$50 per ton. Note that the Japanese government offered a variety of tax and subsidy incentives since 2009. Our demand estimation accounts for these, and hence, this incentive scheme is implemented on top of these existing incentives.

Counterfactual 3. Eco-car Sharing

We implement this scenario as the introduction of a new ownership option by the government. Specifically, the government supplies, perfectly elastically, a platform for *n*-household sharing of ownership of an eco-friendly car. That is, the consumer pays for the rental price on the shared vehicle itself, but the fixed cost of establishing and maintaining the car-sharing platform is paid by the government. Hence, from the consumer's eyes, the supply of this option is perfectly elastic. Then, the consumer faces the essential trade-off between the shared cost of vehicle ownership/utilization versus the reduced rate of utilization. The question is, who 'buys' this option, where and how much?

Our structural approach can get at this question. *n*-household sharing enables consumers to co-own cars, reducing the rental cost of vehicle ownership by 1/n. However, *n*-household sharing causes some inconvenience to the sharing users, reducing the rate of utilization as much as 1/n.¹⁶ Given our empirical model in **Section 3**, we can simulate the consumer's indirect utility of an *n*-household car-sharing option as

$$u_{ijs} = \rho \ln\left(y_i - \frac{1}{n}r_{ij}\right) + \frac{1}{n}v_{ijs}(X_i, Z_j, S_{is}) + \epsilon_{ijs}.$$
(5)

In implementation, we simply add this new option to the set of alternatives while varying

¹⁶There is some uncertainty as to the utility loss from car sharing. On one hand, the utilization loss may not be as much as 1/n. A majority of consumers use vehicles for limited times of a day and a week. By scheduling the times of use, n users (households) may be able to satisfy all driving needs. This may be particularly true if the information and autonomous vehicle technologies allow consumers to costlessly schedule their use in the future. On the other hand, there may be some mental cost associate with scheduling friction, which may be added to the loss of utility. This utility loss can be large when the size of platform users is small, particularly in non-urban areas, but may become negligible if the size of users get large (the economy of agglomeration). Hence, to be conservative, we assume the maximum utility loss of 1/n for this simulation scenario.

the size of n. From (5), it is clear that a consumer who has a low valuation for the benefit of vehicle utilization relative to the value of income tends to buy this option. Hence, consumer behavior under this counterfactual depends on the estimates of the marginal utility parameters with respect to these two terms.

There is a remaining question as to what type of vehicles should be promoted as the sharing platform. We focus on a hybrid vehicle. Hybrid cars offer substantially better mileage per liter of gasoline than the keicars, holding car displacement, size, weight and other on-vehicle amenities. Because consumers in low density areas rely on cars for daily commuting and transportation, allowing those consumers to use hybrid cars instead of keicars is likely to reduce pollution and increase welfare. Nonetheless, we have seen that the consumers tend to own keicars instead of hybrid vehicles and that the estimated demand elasticities can explain this puzzle. Then an interesting question is, Is the sharing of a hybrid vehicle effective in reducing vehicle-related CO_2 emissions?

7.2. Why Does Spatial Heterogeneity Matter?

Before turning to the results, we provide a brief overview of why and how our model with the geographically-explicit portfolio effect may generate different economic impacts, both qualitatively and quantitatively, from previous models.

First, the presence of preferences for a particular clean car can increase the valuation of all ownership portfolios that come with that clean car, making the demand for clean cars 'sticky' and less elastic with respect to the price increase. Here, there are reasons why the portfolio effect may pose some threat to a traditional price instrument through a 'withinportfolio' substitution effect. That is, a consumer may absorb the (relative) price increase of dirty cars by reallocating her spending between the cars in her preferred portfolio, instead of either switching to clean cars or decreasing the number of her vehicle holdings. Second, this portfolio effect may interact, in an important way, with public transit density and other geographic attributes. For example, if households in low density areas value fuel economy more and also value the clean-car portfolio more, then we may expect a larger 'within' substitution effect, but a smaller 'between' substitution effect (including the substitution to the outside option) in low density areas than in high density areas. This implies a higher economic incidence (or a larger welfare loss) in low density areas than in high density areas and that the portfolio effect tends to amplify that effect. Lastly, the portfolio effect can also interact with geographic distribution of income. Recall that we also take into account demand heterogeneity due to differences in income as well as access to public transit, and that consumers have more price-elastic demand for hybrid cars than keicars and in low density areas than in high density areas. This is precisely the intricate consequence of spatial distribution of public transit and income. On one hand, average income is lower in low density areas, which makes the demand more elastic *ceteris paribus*. On the other hand, public transit density is lower in low density areas, which makes the demand less elastic *ceteris paribus*. Given the presence of the portfolio effect, which makes the demand for a certain type of cars less elastic, it becomes highly ambiguous as to how the demand response varies over geographic space, for example, when the price of a hybrid car falls relative to the price of a keicar or when the car-sharing option becomes available. This makes it a good question to investigate empirically.

7.3. Results of Counterfactuals

Figure 4 presents our simulation results visually. In this figure, we display percentage changes in (A) car ownership, (B) car utilization, and (C) vehicle CO_2 emissions under three counterfactual policy scenarios relative to the no-policy benchmark. For each of these outcomes, the means and 95-percent confidence intervals are reported against population density deciles instead of public transit density deciles. We do this because our interest lies in understanding how incomes, public transit, and portfolio preferences interact to generate intricate spatial heterogeneity over geographic space, a simple measure of which we use is population density.

The figure signifies three primary findings. First, as expected, carbon tax is more effective in reducing CO_2 emissions than the other two policies, but its impact has a non-linear relationship to the population density. The CO₂-reducing effect of carbon tax is generally larger in urban areas than in non-urban areas, but eventually flattens out — i.e., the percentage reduction (relative to no policy) becomes invariant with respect to the population density. This occurs precisely due to the intricate interaction between income, public transit, and portfolio preferences. In urban areas, car owners have higher incomes, own a small number of relatively expensive cars, and hence, have inelastic demand for car ownership. The carbon tax, therefore, is not as effective in urban areas in spite of the abundant public transit availability. In moderately dense areas (i.e., suburban areas), public transit is still relatively abundant, yet consumers still use cars for daily errands. The carbon tax, which increases the operating cost of car utilization, is quite effective in reducing both car ownership and utilization in such areas. As the public transit becomes less available, however, consumers have no other means but to use cars for daily commuting and errands, and hence, their demand for car ownership is also inelastic, despite their lower incomes on average. Recall our discussion that households in low density areas tend to own keicars (see Section 2), which are less elastic than other car types (see Section 6). As a result, the effect of carbon tax flattens out.

Second, somewhat unexpectedly, the eco-car sharing policy has a sizable impact on CO_2 emissions, roughly comparable to the carbon tax, but its impact has a U-shaped relationship to the population density. Recall that we assume away supply side responses, and hence, we can explain this result solely by demand side responses. In **Panel A**, we see that the policy indeed has a larger effect on car ownership than the carbon tax, and has the largest impact in moderately dense areas. This occurs mainly because the demand for eco-car sharing is higher in urban areas where the car ownership demand is low (see **Table 5**). In urban areas, both car owners and non-owners switch to the car-sharing option, and hence, the overall impact on car ownership is small. In non-urban areas, car owners switch to the car-sharing option, but the car-sharing demand is weak, and thus, the overall impact is small. In moderately dense areas, the impact is large because the demand for car-sharing is relatively strong and is able to induce car owners rather than non-owners into that option. Interestingly, from **Panel B**, we see a sizable 'rebound effect' at play in urban areas. That is, the sharing of hybrid vehicles lowers both the cost of car ownership and the cost of car utilization. The former induces non-car owners into driving some non-zero distance while the latter induces car-share users into driving more. **Panel B** suggests that the effects are strong in urban areas. **Panel C** indicates, however, that the effect on car ownership is stronger than that on car utilization, and thus, the combined impact translates into a sizable reduction in CO_2 emissions.

Third, feebates are highly *ineffective* not only in reducing CO_2 emissions, but also in inducing a shift toward fuel efficient vehicles. This may come as a surprise, but is indeed consistent with earlier empirical studies. In our stimulation, feebates are implemented in a theoretically justifiable manner, with the SCC value of \$50. As a result, the feebates in our simulation range from -\$35 to \$76, which are added to an annualized rental price, at the maximum. These monetary incentives are far smaller than those implemented in practice.

We now examine the effects of counterfactual policies on substitution patterns in **Table 5**. For ease of interpretation, we report the counterfactual outcomes in average per household for each subsample sorted by population density quintile as of the 2015 level. Note that all statistics reported in this table are *unconditional* averages; this is in contrast to **Table 1** where the descriptive statistics are reported as *conditional* averages (i.e., of those who own cars only). Furthermore, our demand estimation is restricted to the holding of at most two cars. Hence, the counterfactual results are directly comparable to the top panel of **Table 5**, but not to **Table 1**.

Recall our discussion in **Subsection 7.2** that our model is expected to generate geographydependent substitution patterns *within* as well as *between* portfolios, and hence, we also expect policy impacts to exhibit such differential substitution patterns. **Table 5** indeed confirms such expectation. First, carbon tax induces a larger decrease in the ownership share of *keicars* and a larger increase in the share of hybrid cars in non-urban areas than feebates, despite the fact that *both* policies tend to increase relative utility of fuel efficient cars. This occurs because the two policies affect two different price margins. Carbon tax changes the relative operating costs of cars while feebates change the relative ownership costs of cars directly. Given our parameter estimates, the latter induces consumers to substitute away from keicars to hybrid cars in non-urban areas (where demand for car utilization is high) whereas the former causes consumers to shift uniformly toward fuel efficient cars. Second, the share of those who opt for shared eco-cars is monotonically increasing with population density. This is consistent with our discussion in **Subsection 7.1**. From eq. (5), a consumer who has a low valuation for the benefit of vehicle utilization relative to the value of income tends to buy this option. The share of such consumers tend to monotonically increase with the population density, and hence, the result. Third, feebates are not effective in reducing car ownership or utilization, but are effective in inducing a shift toward fuel efficient vehicles. Both are expected in theory and echo previous empirical findings.

Lastly, we touch on consumer welfare. Economic theory predicts that consumers with inelastic demand would have higher tax incidence, and thus, lose more from a given level of tax while consumers would gain more from a new product when their demand for the product is high. Thus, given our results above, we expect the consumers in low-density areas would lose more from the carbon tax than in high-density areas while the consumers in high-density areas would gain more from the eco-car sharing policy. For feebates, the effect of spatial heterogeneity is, a priori, unclear. Consumers, either with or without inelastic demand, would be able to switch from less fuel efficient cars, which are taxed more under feebates, to more fuel efficient cars, which are subsidized. Hence, those who have inelastic demand for gas-guzzlers would lose more from the policy. We confirm these points in **Figure 5**, where we plot changes in compensating variation by population density. To estimate compensating variation, we use the standard formula à la Small and Rosen (1981). A challenge, however, is that our specification involves a non-linear income term, and thus, we do not have a closed-form analytical expression for the compensating variation term. We thus follow Herriges and Kling (1999), and solve for it numerically.

As shown in **Figure 5**, the welfare loss from the carbon tax policy varies substantially over geographic space, ranging 100,000 yen per year per household in lowest-density areas to less than 10,000 yen in highest-density areas. The overall welfare loss averages at roughly 59,7000 yen per year per household. This estimate is roughly comparable to those in Bento et al. (2009), who estimate that the average welfare loss is around 22,000 yen for a 6.6 yen gasoline tax increase when there is no revenue recycling. In contrast, the welfare gain from the eco-car sharing policy range from 100,000 yen in highest-density areas to 25,000 yen in lowest-density areas, with the average welfare gain of roughly 58,600 yen per year per household. For feebates, the welfare effect is indistinguishably zero, and does not vary much over space. Some consumers continue to own the same portfolio of vehicles while others switch to a new portfolio. Some of the switchers gain while others lose; some of the stayers gain while others lose. The losers and gainers are roughly balanced because we simulate the feebates in a theoretically consistent manner — fees and rebates are distributed equally over car models.

We emphasize again, however, that these welfare responses only account for demand side only, and thus, neither supply side responses (i.e., price/quality change) nor changes in producer surplus are taken into account. Nonetheless, this contrast between the three policies (in terms of welfare) suggest a potential for a policy mix, a carbon tax combined with incentives for eco-car sharing, which in our view, has not been fully explored to date.

8. Conclusion

How does demand for vehicle transport vary over geographic space? How does the economic impact of a policy to control vehicle emissions depend on such spatial heterogeneity? With these questions in mind, we estimate a model of vehicle ownership and utilization, explicitly accounting for the role of incomes, public transit networks, and portfolio considerations, on a large cross-sectional sample of households in Japan.

Our model provides evidence that consumers in low density areas inelastically demand keicars (i.e., extremely small vehicles) for its cost performance and that all three factors — incomes, public transit, and portfolio considerations — are important in generating this inelastic demand. Low-income households in low density areas tend to have more elastic demand than high-income households in high density areas, yet their demand for keicars is more inelastic than for hybrid cars because they need low-cost fuel-efficient vehicles for daily use and because they opt for a combination of keicars with others. Hence, the estimated model can resolve two empirical puzzles we observe in Japan: Overall price elasticity of demand for vehicle ownership increases with vehicle size; invariance of demand for hybrid vehicles with respect to public transit density.

To demonstrate the economic significance of these findings, we use the estimated demand to simulate the effects of three counterfactual climate mitigation policies: carbon tax, feebates, and eco-car sharing. Our simulation results indeed confirm highly geographydependent substitution patterns *within* as well as *between* portfolios, and hence, policy impacts exhibit a highly non-linear relationship to population density. Car owners in urban areas are generally richer, prefer having a smaller number of cars, and thus, have relatively inelastic demand for car ownership than those in non-urban areas. Car owners in non-urban settings, on the other hand, have inelastic demand for car utilization and have preferences for a mix of keicars with other cars, whose demand is estimated to be highly inelastic, precisely due to low public transit availability. We obtain three important findings as a results of the mixed effects of these. First, carbon tax has a larger CO_2 -reducing impact in non-urban settings, yet the effect flattens out as the transit density further declines as the vehicle demand becomes increasingly *sticky*. Consequently, the welfare loss from carbon tax is greatest in the lowest density areas. Second, eco-car sharing can mitigate this welfare penalty while remaining equally effective. Third, feebates perform poorly relative to either policy. Taken all together, our results suggest the potential for a welfare-improving policy mix: a carbon tax in combination with a car-sharing platform, particularly in non-urban settings.

There are several important limitations to our study. First, we simulate the eco-car sharing policy as an introduction of a purely hypothetical new ownership option. Hence, the counterfactual policy effect depends heavily on the way the new option is modeled (as well as estimated parameters, of course). Second, our analysis is limited to the demand side only. Hence, neither supply-side responses (e.g., price/quality changes) nor producer surplus is accounted for in our simulation. These issues can be addressed only with availability of suitable data, and thus, are left for future research. Nonetheless, our work has important implications for climate mitigation policies and for a new direction of research. In the era of the rising gig economy and autonomous vehicle technology, sharing of vehicles may become increasingly costless and seamless, both at the ownership level or at the utilization level. On one hand, the supply of car-sharing service is likely to enjoy the economy of density. Hence, the economic cost of such service is likely to low in high density areas. On the other hand, the social demand for such service (i.e., the sum of private and external benefits) may be high in low density areas. If so, such a mismatch between demand and supply would suggest the need for the government intervention. Our results shed some lights on this conjecture — unexploited demand for sharing of a eco-friendly vehicle is large, it has the potential for reducing carbon emissions on a scale comparable to the carbon tax, and it may work as the means to counteract the welfare-decreasing effect of the carbon tax, particularly in non-urban settings.

References

- [1] Allcott, Hunt and Nathan Wozny. 2014. Gasoline Prices, Fuel Economy, and the Energy Paradox. *Review of Economics and Statistics* 96(5): 779–795.
- [2] Anderson, Soren T., Ian W. H. Parry, James M. Sallee, and Carolyn Fischer. 2011. Automobile Fuel Economy Standards: Impacts, Efficiency, and Alternatives. *Review of Environmental Economics and Policy* 5(1): 89–108.
- [3] Andersson, Julius J. 2019. Carbon Taxes and CO2 Emissions: Sweden as a Case Study. American Economic Journal: Economic Policy 11(4): 1–30.
- [4] Beaudoin, Justin and Cynthia Lin Lawell. 2018. The Effects of Public Transit Supply on the Demand for Automobile Travel. Journal of Environmental Economics and Management 88: 447–467.
- [5] Beresteanu, Arie and Shanjun Li. 2011. Gasoline Prices, Government Support, and the Demand for Hybrid Vehicles in the United States. *International Economic Review* 52(1): 161–182.
- [6] Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. Automobile Prices in Market Equilibrium. *Econometrica* 63(4): 841–890.
- [7] —. 1999. Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy. American Economic Review 89 (3): 400–430.
- [8] Bento, Antonio, Lawrence Goulder, Mark Jacobsen, and Roger von Haefen. 2009. Distributional and Efficiency Impacts of Increased US Gasoline Taxes. American Economic Review 99(3): 667–699.
- [9] Bento, Antonio, Maureen Cropper, Ahmed Mobarak, and Katja Vinha. 2005. The Effects of Urban Spatial Structure on Travel Demand in the United States. *Review of Economics and Statistics* 87(3): 466–478.
- [10] Boarnet, Marlon and Randall Crane. 2001. The Influence of Land Use on Travel Behavior: Specification and Estimation Strategies. *Transportation Research Part A: Policy* and Practice 35(9): 823–845.
- [11] Boyle, Kevin J., Mark Morrison, Darla Hatton MacDonald, Roderick Duncan, and John Rose. 2016. Investigating Internet and Mail Implementation of Stated-Preference Surveys While Controlling for Differences in Sample Frames. *Environmental and Resource Economics* 64: 401–419.

- [12] Brownstone, David and Kenneth Train. 1998. Forecasting new product penetration with flexible substitution patterns. *Journal of Econometrics* 89(1-2), 109–129.
- [13] Busse, Meghan R., Christopher R. Knittel, and Florian Zettelmeyer. 2013. Are Consumers Myopic? Evidence from New and Used Car Purchases. *American Economic Review* 103(1): 220–256.
- [14] Dahl, Gordon. 2002. Mobility and the Return to Education: Testing A Roy Model with Multiple Markets. *Econometrica* 70(6): 2367–2420.
- [15] D'Haultfœuille, Xavier, Pauline Givord, and Xavier Boutin. 2014. The Environmental Effect of Green Taxation: The Case of the French Bonus/Malus. *The Economic Journal* 124(578): 444–480.
- [16] Dorband, Ira Irina, Michael Jakob, Matthias Kalkuhl, and Jan Christoph Steckel. 2019.
 Poverty and distributional effects of carbon pricing in low- and middle-income countries
 A global comparative analysis. World Development 115: 246–257.
- [17] Domencich, Tom and Daniel L. McFadden. 1975 (Reprinted 1996). Urban Travel Demand: A Behavioral Analysis. North-Holland Publishing Co.
- [18] Dubin, Jeffrey and Daniel McFadden. 1984. An Econometric Analysis of Residential Electric Appliance Holdings and Consumption. *Econometrica* 52(2): 345–362.
- [19] Duranton, Gilles and Matthew Turner. 2011. The Fundamental Law of Road Congestion: Evidence from US Cities. American Economic Review 101(6): 2616–2652.
- [20] Fremstad, Anders and Mark Paul. 2019. The Impact of a Carbon Tax on Inequality, *Ecological Economics* 163: 88–97.
- [21] Gentzkow, Matthew. 2007. Valuing New Goods in a Model with Complementarity: Online Newspapers. American Economic Review 97(3): 713–744.
- [22] Gillingham, Kenneth. 2014. Identifying the elasticity of driving: Evidence from a gasoline price shock in California, *Regional Science and Urban Economics* 47: 13–24.
- [23] Gillingham, Kenneth, Alan Jenn, and Inês M.L. Azevedo. 2015. Heterogeneity in the response to gasoline prices: Evidence from Pennsylvania and implications for the rebound effect, *Energy Economics* 52(1): S41–S52.
- [24] Gillingham, Kenneth, Matthew J. Kotchen, David S. Rapson, and Gernot Wagner. 2013. The rebound effect is overplayed. *Nature* 493, 475–476.

- [25] Glaeser, Edward and Matthew Kahn. 2010. The Greenness of Cities: Carbon Dioxide Emissions and Urban Development. *Journal of Urban Economics* 67(3): 404–418.
- [26] Goldberg, Pinelopi Koujianou. 1998. The Effects of the Corporate Average Fuel Efficiency Standards in the US. Journal of Industrial Economics 46(1): 1–33.
- [27] Goulder, Lawrence H., Marc A.C. Hafstead, GyuRim Kim, and Xianling Long. 2019. Impacts of a carbon tax across US household income groups: What are the equityefficiency trade-offs? *Journal of Public Economics* 175: 44–64.
- [28] Grainger, Corbett A. and Charles D. Kolstad. 2010. Who Pays a Price on Carbon? Environmental and Resource Economics 46: 359–376.
- [29] Herriges, Joseph A. and Catherine L. Kling. 1999. Nonlinear Income Effects in Random Utility Models. *Review of Economics and Statistics* 81(1): 62–72.
- [30] Klier, Thomas and Joshua Linn. 2015. Using Taxes to Reduce Carbon Dioxide Emissions Rates of New Passenger Vehicles: Evidence from France, Germany, and Sweden. *American Economic Journal: Economic Policy* 7(1): 212–242.
- [31] Knittel, Christopher R. 2011. Automobiles on Steroids: Product Attribute Trade-Offs and Technological Progress in the Automobile Sector. American Economic Review 101(7): 3368–3399.
- [32] —. 2012. Reducing Petroleum Consumption from Transportation. Journal of Economic Perspectives 26(1): 93–118.
- [33] Konishi, Yoshifumi and Meng Zhao. 2017. Can Green Car Taxes Restore Efficiency? Evidence from the Japanese New Car Market. Journal of the Association of Environmental and Resource Economists 4(1): 51–87.
- [34] Levinson, David M. and Ajay Kumar. 1997. Density and the Journey to Work. Growth and Change 28(2):147–172.
- [35] Lindhjem, Henrik and Ståle Navrud. 2011. Are Internet surveys an alternative to faceto-face interviews in contingent valuation? *Ecological Economics* 70(9): 1628–1637.
- [36] Maestre-Andres, Sara, Stefan Drews, and Jeroen van den Bergh. 2019. Perceived fairness and public acceptability of carbon pricing: a review of the literature. *Climate Policy* 19(9): 1186–1204.

- [37] McFadden, Daniel and Kenneth Train. 2000. Mixed MNL models for discrete response. Journal of Applied Econometrics 15: 447–470.
- [38] Nielsen, Jytte Seested. 2011. Use of the Internet for willingness-to-pay surveys: A comparison of face-to-face and web-based interviews. *Resource and Energy Economics* 33(1): 119–129.
- [39] Revesz, Richard, Michael Greenstone, Michael Hanemann, Michael Livermore, Thomas Sterner, Denise Grab, Peter Howard, and Jason Schwartz. 2017. Best Cost Estimate of Greenhouse Gases. *Science* 357(6352), 655.
- [40] Small, Kenneth A. and Harvey S. Rosen. 1981. Applied Welfare Economics with Discrete Choice Models. *Econometrica* 49(1): 105–130.
- [41] Terza, Joseph, Anirban Basu, and Paul Rathouz. 2008. Two-stage Residual Inclusion Estimation: Addressing Endogeneity in Health Econometric Modeling. *Journal of Health Economics* 27(3): 531–543.
- [42] Train, Kenneth. 1986. Qualitative Choice Analysis: Theory, Econometrics, and an Application to Automobile Demand. Cambridge: MIT Press.
- [43] Wakamori, Naoki. 2015. Portfolio Considerations in Differentiated Product Purchases: An Application to the Japanese Automobile Market. SFB/TR Discussion Paper 499.
- [44] West, Sarah E. 2004. Distributional Effects of Alternative Vehicle Pollution Control Policies. Journal of Public Economics 88: 735–757.
- [45] Wooldridge, Jeffrey. 2015. Control Function Methods in Applied Econometrics. Journal of Human Resources 50(2): 420–445.



ø.

Public transit density

Ņ

0

0

5000

Figure 1. Measures of Access to Public Transit

Note: All observations are at the municipality level. For reference only, we use county- or city-level population density data for U.S. cities.

• Rail

20000

10000 15000 Population density (per km2) + Bus

25000

3

0

0

2

.4 .6 Public transit density in 1980 Observed Frac. poly. fit

.8

45 degree line



Note: For the number of cars owned, the mean and the 80th percentile are reported; Vehicle kilometer traveled (VKT) is the sum over two most frequently used cars owned by each household; Fuel economy ratings and the share of hybrid cars/*keicars* are the (unweighted) averages over the two cars; Household income is in millions of Japanese yen; For vehicle portfolios, we report (i) the share of complementary portfolios, i.e., the share of households who own any combination of *different* vehicle types (keicar, regular, and minivan), and (ii) the share of households who own a keicar with other vehicle type, only among those who own two cars. All figures, with the exception of panel A, report the averages of these indicators over households in each public transit density bin.



Figure 3. Prediction Performance: Full Model versus Naïve Random Coefficient (RC) Logit

Note: Predicted shares in this figure are calculated as *unconditional* shares: i.e., not conditional on car holding unlike in Figure 2.



Figure 4. Heterogeneous Impacts of Counterfactual Policies over Geographic Space

Note: The vertical axis is the percentage change under each counterfactual relative to the no-policy benchmark. For carbon tax and feebates, we use the SCC value of $50/\text{ton-CO}_2$. For eco-car sharing, we set the number of households sharing a hybrid vehicle at 3. The whisker indicates the 95% confidence interval for bootstrap iterations.



Figure 5. Welfare Impacts of Counterfactual Policies over Geographic Space

Note: The vertical axis denotes changes in compensating variation in 10,000 yen (at \$1 = 100 yen) per year per household relative to the no-policy benchmark. For carbon tax and feebates, we use the SCC value of $\$50/ton-CO_2$. For eco-car sharing, we set the number of households sharing a hybrid vehicle at 3. The whisker indicates the 95% confidence interval for bootstrap iterations.

				PO	pulation Del	isity Quinti	es			
	Lowe	est	4tł	1	3rc	1	2no	d	Highe	est
Number of obs. (households)	20,9	63	20,8	51	21,04	41	21,0	78	21,03	33
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Public transit density	0.112	0.064	0.178	0.080	0.277	0.125	0.392	0.110	0.559	0.119
Congestion measure	0.595	0.108	0.689	0.105	0.753	0.136	0.813	0.103	0.899	0.062
Household income (10,000 yen)	601.6	410.1	634.6	427.1	656.1	442.5	684.1	460.7	746.1	530.1
Household size [Num. of people in household]	2.94	1.39	2.88	1.32	2.81	1.27	2.70	1.24	2.52	1.25
Own car? [1 = yes]	0.938	0.241	0.908	0.289	0.829	0.376	0.730	0.444	0.497	0.500
Drive to work? [1 = yes]	0.548	0.498	0.434	0.496	0.266	0.442	0.154	0.361	0.070	0.256
Num. of cars owned [of those who own car]	1.73	0.90	1.56	0.78	1.32	0.61	1.17	0.44	1.09	0.36
<i>Most frequently used car</i> Monthly VKT	1,162.7	2,378.0	1,000.6	1,916.8	892.3	1,884.1	785.6	1,837.6	730.3	1,785.4
Price (10,000 yen)	199.6	129.3	206.5	135.4	218.9	140.0	229.7	150.7	261.7	182.1
Fuel economy (km/L)	20.4	7.6	20.5	7.7	19.8	7.6	19.3	7.7	18.5	7.7
Vehicle size (mm) [length + width + height]	7,276.2	819.0	7,297.1	801.1	7,394.9	783.4	7,457.4	793.8	7,537.6	844.1
Hybrid [1 = yes]	0.135	0.342	0.146	0.353	0.146	0.353	0.150	0.357	0.150	0.357
Keicar [1 = yes]	0.349	0.477	0.327	0.469	0.289	0.454	0.254	0.436	0.221	0.415
Second most used car										
Monthly VKT	964.4	1,920.3	940.7	2,231.6	937.0	2,318.8	941.9	2,856.1	1,065.5	3,399.3
Price (10,000 yen)	170.8	116.4	180.8	128.6	194.7	151.9	215.9	184.1	283.5	268.5
Fuel economy (km/L)	21.1	7.0	21.0	7.1	21.0	7.3	20.2	7.6	19.0	7.3
Vehicle size (mm) [length + width + height]	7,037.7	760.9	7,058.4	777.2	7,068.4	838.0	7,192.5	843.2	7,234.9	1,050.5
Hybrid [1 = yes]	0.039	0.195	0.035	0.183	0.023	0.150	0.015	0.122	0.009	0.096
Keicar [1 = yes]	0.531	0.499	0.501	0.500	0.480	0.500	0.420	0.494	0.329	0.470
<i>Portfolio shares</i> [of those who own two (or more	e) cars]									
Keicar-Keicar	12.4%		10.8%		11.5%		10.5%		7.2%	
Keicar-Regular	26.9%		25.9%		25.4%		25.4%		21.4%	
Keicar-Minivan	32.4%		31.8%		31.7%		31.2%		29.5%	
Regular-Regular	6.8%		8.7%		9.0%		9.4%		11.9%	
Regular-Minivan	12.6%		12.9%		14.0%		13.4%		16.3%	
Minivan-Minivan	8.9%		9.9%		8.4%		10.2%		13.7%	
Total share	100.0%		100.0%		100.0%		100.0%		100.0%	

Table 1. Descriptive Statistics by Population Density Quintiles

Population Density Quintiles

		No Portfolio Effe	cts	Π	. Portfolio Effect	S	III. Po	ortfolio Effects +	2SRI
	M	Interacti	on Terms	Mood	Interacti	on Terms	Moon	Interactio	on Terms
	Mean Parameter	Transit Density	Household Size	Parameter	Transit Density	Household Size	Parameter	T ransit Density	Household Size
ln(y - r)	5.495 *** 001070			5.512 *** 0 107)			5.570 *** (0 108)		
ХРК	-0.295 ***	0.1575 ***	0.082 ***	-0.323 ***	0.187 ***	0.086 ***	-0.331 ***	0.206 ***	0.086 ***
HP /W/	(0.008) 39.017 ***	(0.015) 9 942 ***	(0.002) -11.929 ***	(0.009) 46.350 ***	(0.016) 8.873 ***	(0.002) -13 932 ***	(0.009) 46 509 ***	(0.017) 8.101 ***	(0.002) -13 969 ***
	(1.251)	(2.313)	(0.279)	(1.343)	(2.482)	(0.326)	(1.371)	(2.630)	(0.326)
Size	-0.424 *** (0.010)	-0.832 *** (0.018)	0.086 *** (0.002)	-0.487 *** (0.011)	-0.853 *** (0.020)	0.112 *** (0.003)	-0.463 *** (0.012)	-0.905 *** (0.021)	0.113 *** (0.003)
Portfolio Effects									
Kei-Kei				-0.063	-0.547 **	-0.164 ***	-0.179 *	-0.314	-0.164 ***
				(0.095)	(0.254)	(0.020)	(0.095)	(0.255)	(0.020)
Kei-Regular				0.329 ***	-0.673 ***	-0.165 ***	0.205 ***	-0.408 **	-0.164 ***
Kei-Minivan				(coo.o) *** 696.0	-0.553 ***	-0.249 ***	0.855 ***	-0.320 *	-0.249 ***
Remilar-Minivan				(0.073) 1 040 ***	(0.172) _1 738 ***	(0.016) _0 191 ***	(0.073) 0.0218 ***	(0.174) -0 999 ***	(0.016) _0 191 ***
itegual Trittin att				(0.098)	(0.242)	(0.023)	(0.099)	(0.243)	(0.023)
Minivan-Minivan				1.290 *** (0.386)	-2.109 ** (0.973)	-0.308 *** (0.075)	1.1846 *** (0.386)	-1.907 * (0.975)	-0.308 *** (0.075)
Used-car dummy		Yes			Yes			Yes	
x Metropolitan dummies									
Make dumnies		Yes			Yes			Yes	
x Metropolitan aummes Endlatine diimmise		Vac			Vac			Vec	
x Metropolitan dummies		51			103			5	
Vehicle-type dummies x Metropolitan dummies		Yes			Yes			Yes	
# of 0bs.		13,890,396			13,890,396			13,890,396	
# of Cases		82,730			82,730			82,730	

Table 2. Vehicle Ownership Demand

	(1)	(2)	(3)	(4)	(5)	(9)	(7)	(8)
ln(y - r)	0.016	0.043 *	0.023	0.043 ** (ccn.o.)	0.075 **	0.100 **	0.078 **	0.114 ***
imes Transit Density	(010.0) ***	-0.119 ***	-0.104 ***	-0.121 ***	(ccu.u) -0.082 *	(*cu.u) ** 660.0-	(0c0.0) -0.089 **	-0.103 **
\times Congestion	(0.040)	(0.042)	(0.040)	(0.039)	(0.046) -0.082 *	(0.046) -0.082	(0.044) -0.079	(0.049) -0.098 *
ln(YPK)	-0.339 ***	-0.361 ***	-0.340 ***	-0.363 ***	(0.049) -0.391 ***	(0.050) -0.455 ***	(0.050) -0.390 ***	(0.051) -0.447 ***
imes Transit Density	(0.078) 0.276 **	(0.078) 0.336 **	(0.080) 0.274 **	(0.075) 0.335 ***	(0.123) 0.224	(0.121) 0.267 *	(0.122) 0.222	(0.127) 0.268 *
\times Congestion	(0.128)	(0.132)	(0.129)	(0.118)	(0.152) 0.092	(0.146) 0.155	(0.144) 0.088	(0.157) 0.138
HP/W	0.497	0.480	0.526	0.493	(0.158) 0.558 (1.220)	(0.161) 0.503 (1.200)	(0.160) 0.553 (1.250)	(0.163) 0.513
Size	0.274 ***	0.282 ***	0.274 *** 0.274 ***	0.281 *** 0.281 ***	(ссст) 0.272 ***	0.281 ***	0.273 ***	0.281 *** 0.281 ***
Second Car	(0.042) -0.211 *** (0.033)	(0.043) -0.211 *** (0.032)	(0.044) -0.211 *** (0.034)	(0.042) -0.211 *** (0.031)	(0.044) -0.210 *** (0.032)	(0.045) -0.211 *** (0.033)	(0.044) -0.211 *** (0.032)	(0.044) -0.211 *** (0.032)
Demographic controls Urban structure controls Metropolitan dummies Selection controls	>>	>> >	```	>>>>	>>	>> >	>>>	›››
χ^2 -stat. on selection terms	I	21.21 ***	ł	27.39 ***	I	38.14 ***	;	31.62 ***
Num. of obs. R ²	21,456 0.075	21,456 0.078	25,953 0.077	21,456 0.078	21,456 0.076	21,456 0.078	25,953 0.077	21,456 0.078

Table 3. Vehicle Utilization Demand

Note: Bootrsrap standard errors in parenthesis (500 replications).

		Public	transit quintiles		
	5th (Lowest)	4th	3rd	2nd	1st (Highest)
Car ownership elastic	city				
Hybrid cars	-1.009 ***	-0.892 ***	-0.706 ***	-0.483 ***	-0.296 ***
	(0.022)	(0.018)	(0.011)	(0.008)	(0.006)
Diesel cars	-0.867 ***	-0.825 ***	-0.704 ***	-0.546 ***	-0.365 ***
	(0.027)	(0.020)	(0.015)	(0.017)	(0.013)
Kei-cars	-0.749 ***	-0.643 ***	-0.488 ***	-0.329 ***	-0.223 ***
	(0.016)	(0.014)	(0.010)	(0.007)	(0.005)
VKT elasticity					
w.r.t. Net income	0.044 **	0.030	0.010	-0.007	-0.033
	(0.021)	(0.020)	(0.018)	(0.017)	(0.019)
w.r.t. YPK	-0.338 ***	-0.306 ***	-0.267 ***	-0.224 ***	-0.167 **
	(0.073)	(0.072)	(0.070)	(0.069)	(0.078)

Table 4. Elasticity Estimates

Note: Bootstrapped standard errors in parenthesis. We use 50 replications for car ownership elasticity and 500 replications for VKT elasticity.

	Popu	lation dens	ity quintile	s (as of 20	15)	
	5th (Lowest)	4th	3rd	2nd	1st (Highest)	All
No Policy Counterfactual						
Pop. density (1000/km2)	0.273	1.082	3.409	7.414	14.412	5.318
Pub. transit density	0.112	0.178	0.277	0.392	0.559	0.303
Num. of cars owned	1.130	0.995	0.799	0.620	0.426	0.794
Car/ownership type (%)						
Hybrid	7.10	6.97	5.50	4.15	2.56	5.26
Diesel	0.92	0.95	0.87	0.72	0.50	0.79
Keicar	39.52	28.83	20.00	12.70	7.88	21.79
Car-sharing	-	-	-	-	-	-
Monthly VKT (km)	939.16	713.75	447.51	273.83	112.33	497.32
CO ₂ from driving (kg/mo.)	116.94	90.92	58.34	36.79	15.47	63.69
Counterfactual I: Carbon tax (\$	50/ton-CO2)				
Num. of cars owned	1.105	0.972	0.781	0.607	0.420	0.777
Car/ownership type (%)						
Hybrid	7.24	7.08	5.55	4.16	2.55	5.32
Diesel	0.94	0.96	0.87	0.71	0.50	0.80
Keicar	38.75	28.25	19.64	12 51	7.80	21 39
Car-sharing	-	-	-	-	-	-
Monthly VKT (km)	00E 00	601 22	127 20	262.02	100 / 7	171 00
CO from driving (log/month)	110 (5	001.23	427.30	202.02	14.00	(0.20
CU_2 from driving (kg/month)	110.65	80.07	55.32	35.00	14.89	60.39 F 07
Cv (10,000 yen/year)	-9.41	-8.81	-0.13	-3.84	-1.04	-5.97
Counterfactual II: Feebates (me	edian as a ta	x-subsidy	cutoff)			
Num. of cars owned	1.130	0.995	0.799	0.620	0.426	0.794
Car/ownership type (%)						
Hybrid	7.13	6.99	5.52	4.17	2.57	5.28
Diesel	0.93	0.95	0.87	0.72	0.50	0.79
Keicar	39.57	28.87	20.04	12.72	7.89	21.82
Car-sharing	-	-	-	-	-	-
Monthly VKT (km)	938.96	713.54	447.32	273.67	112.25	497.15
CO ₂ from driving (kg/month)	116.80	90.81	58.27	36.74	15.45	63.61
CV (10,000 yen/year)	-0.01	-0.03	-0.03	-0.04	-0.03	-0.03
Counterfactual III: Car-sharing	of hybrid ca	rs (3-hous	sehold sha	ring)		
Num. of cars owned	1.099	0.962	0.770	0.599	0.419	0.770
Car/ownership type (%)						
Hybrid	6.82	6.62	5.15	3.84	2.35	4.95
Diesel	0.89	0.90	0.81	0.66	0.46	0.74
Keicar	37.98	27.40	18.74	11.73	6.22	20.41
Car-sharing	4.09	5.20	6.66	7.82	8.31	6.41
Monthly VKT (km)	926.25	702.56	441.42	272.10	113.85	491.24
CO ₂ trom driving (kg/month)	114.31	88.36	56.44	35.53	15.06	61.94
CV (10,000 yen/year)	2.91	4.10	5.83	7.52	8.93	5.86

Table 5. Impacts of Counterfactual Policies

Note: We assume the share of those who use car-sharing as a substitute for holding a car is zero in the benchmark. CV stands for the compensating variation.



Appendix A. Sample versus Population Distribution by Prefecture

900 800 Survey sample Population 700 600 500 400 300 200 100 0 Tokyo Kanagawa Niigata Toy ama Ishikawa Fukui Yamagata Fukushima Ibaragi Tochigi Gunma Saitama Chiba Hi roshima Yamaguchi Tokushima Kagawa Ehime Kochi Miyazaki Kagoshima Okinawa Miyagi Akita Nagano Gifu Shiga Kyoto Osaka Hyogo Tottori Shimane Saga Iwate Aichi Mie Nara Oita /amanashi Shizuoka Wakayama Okayama Fukuoka Kumamoto Hokkaido Aomori Nagasaki



(c) Household size

4.0

Appendix B. Choice Set and Key Variables

This appendix explains how we define the choice set for the vehicle ownership estimation and key variables used in estimation. We classify each respondent's observed vehicle choice according to its curb weight, car type, fuel type, sales type, and make as in Table B1. We do essentially the same for the first and the second most frequently used cars. This yields a choice set of 168 alternatives (incl. the option of holding no car). Vehicle attributes such as displacement, horsepower, size, and weight are averaged over observations for each choice alternative. Price variables (i.e., rental price and YPK) are similarly averaged over observations for each choice alternatives and then adjusted to contemporaneous values using information on purchase year/month (see below). This type of aggregation is common in studies that use household-level data [see, for example, Bento *et al.* (2009)]. Even with this level of aggregation, Stata's maximum likelihood estimation of the conditional logit model takes roughly 1 hour for each run on a modern computer (10 cores/20 threads, Core i9 CPU, 64GB memory), due partly to our large sample size. Intuitively speaking, this aggregation implies that the consumer in the model makes her choice comparing the choice of her own against 'average' economic values of alternatives.

Table B1. Vehicle Classification

U.S. NHTSA Classification	Car Type	Fuel Type	Sales Type	Make
Mini: Weight $\leq 900 \text{ (kg)}$	Keicar	Diesel	New	Toyota
Light: $900 < \text{Weight} \le 1,150$	Regular	Hybrid	Used	Honda
Compact: $1,150 < \text{Weight} \le 1,350$	Minivan	Gasoline		Other
Medium: $1,350 < \text{Weight} \le 1,600$				
Heavy: Weight $> 1,600$				

Our key variables used in estimation (both the ownership and the utilization equation):

- **Household income**: Annual before-tax incomes are reported with an interval of 1 million yen from 2 million yen up to 10 million yen, and then 10 to 15 million yen, 15 to 20 million yen, and 20 to 30 million yen. We use the mid-point of income interval as a measure of annual income.
- **Rental price**: The survey records the purchase price of each of the two frequently used cars. We convert the purchase price into a rental price using an annual depreciation rate of 10% and annual interest rates, which is allowed to vary by year/month of purchase. We add annualized automobile taxes and tax incentives to this rental price, which are also allowed to vary by year/month of purchase. The rental price is further adjusted for regional inflation rates.
- Household size: We use the raw number of individuals in the household.
- **Yen per kilometer (YPK)**: YPK is the gasoline price divided by the catalog-based fuel economy ratings. We use the gasoline price for the year/month of purchase.

Horsepower/weight (HP/W): HP/W is the horsepower divided by curb weight.

Size: Vehicle size measured as the sum of length, height, and width.

- Garage certificate dummy: In Japan, basically all vehicles must obtain a garage certificate, but in some areas, only keicar do not require a garage certificate. Since a garage certificate is required to prevent on-street parking, areas where it is not required are generally considered to have sufficient parking space, and such areas may have high demand for vehicles in terms of ownership costs and ease of driving.
- Metropolitan dummies: Four dummy variables indicating whether the respondent resides in major metropolitan areas: Kanto, Chukyo, Kinki, or Kitakyushu.

In the vehicle utilization equation, we also use (1) the respondent's age, years of education, work status, marital status as demographic controls and (2) population density and availability of hospitals and public parks at the district level as geographic controls.

Appendix C. Two-stage Residual Inclusion (2SRI) Method

The full specification of our conditional logit model is given by:

$$u_{i(j,k)s} = \rho \ln(y_i - r_{ij} - r_{ik}) + \mathbf{Z}'_j(\lambda_0 + x_i\lambda_1 + tr_s\lambda_2) + \mathbf{Z}'_k(\lambda_0 + x_i\lambda_1 + tr_s\lambda_2) + (\kappa_0 + \kappa_1x_i + \kappa_2tr_s)I(j,k) + \epsilon_{i(j,k)s}$$

where x_i is the household size and tr_s is the rail transit density.

Our concern is the endogeneity of tr_s that arises through consumer's endogenous residential decision. Consumers who reside in low transit areas may have innate preferences for driving; those who reside in high transit areas may have innate distastes for driving. Such innate preferences may correlate with preferences for certain types of cars, and hence, enter the indirect utility as a choice-specific unobserved error $\epsilon_{i(j,k)s}$. The 2SRI method we describe below can parsimoniously address all endogeneity concerns that arise through consumer's endogenous residential decision.

Let **w** be a vector of valid instruments for tr_s such that:

$$tr_s = g(\mathbf{w}) + \nu_s$$
 and $E[\epsilon_{i(i,k)s} | \mathbf{w}] = 0.$

Let us further assume that we can write

$$E[\epsilon_{i(j,k)s}|\nu_s] = \Lambda_{i(j,k)}\nu_s.$$

Note that we are exploiting the fact that the pure source of endogeneity lies at the district level (i.e., innate preferences for choosing district s), yet is correlated with choice-specific unobservables. Wooldridge (2015) shows that in nonlinear models, we can flexibly apply the control function approach and rewrite the indirect utility as:

$$u_{i(j,k)s} = \rho \ln(y_i - r_{ij} - r_{ik}) + \mathbf{Z}'_j(\lambda_0 + x_i\lambda_1 + tr_s\lambda_2) + \mathbf{Z}'_k(\lambda_0 + x_i\lambda_1 + tr_s\lambda_2) + (\kappa_0 + \kappa_1 x_i + \kappa_2 tr_s)I(j,k) + \nu_s + \nu_s^2 + \Lambda_{i(j,k)}\nu_s + \Lambda_{i(j,k)}\nu_s^2.$$

Given this, we implement this 2SRI method as follows. In the first step, we fit the fractional polynomial regression using the current population density and the past rail transit density as of 1980 as instruments.¹ Stata's fractional polynomial routine has chosen the following as the best linear fit:

$$tr_{s,2015} = \theta_0 + \theta_1 tr_{s,1980} + \theta_2 pop_{s,2015}^{(0.7)} + \theta_3 pop_{s,2015}^{(2.0)} + \nu_s.$$

We then obtain the estimates of the residual $\hat{\nu}_s$. Note here that we *do not* include choice-specific or household-specific covariates in this first-stage regression because the residual of interest varies only at the district level *s*. In the second step, we estimate the conditional logit using $\hat{\nu}_s$, $\hat{\nu}_s^2$, $\Lambda_{i(j,k)}\hat{\nu}_s$ and $\Lambda_{i(j,k)}\hat{\nu}_s^2$ along with other covariates in the original utility. Note, however, that the first two terms $\hat{\nu}_s$ and $\hat{\nu}_s^2$ are canceled out since they are constant across choice alternatives. There is one remaining issue. We have no *a priori* knowledge on $\Lambda_{i(j,k)}$, and without it, we need to estimate them as choice-specific parameters. This results in estimating 168 additional parameters. Our earlier attempt has resulted in non-convergence. Hence, we have chosen an alternative route. By assumption, choice-specific attributes such as \mathbf{Z}_j , \mathbf{Z}_k , and I(j,k) are pre-determined, and thus, their interactions with $\hat{\nu}_s$ should be able to parsimoniously handle this term. Hence, we estimate the second-stage model, interacting $\hat{\nu}_s$ and $\hat{\nu}_s^2$ with these observed choice-specific attributes.

¹We make sure that all of the households included in the estimation do not own cars sold prior to 1980.