Minnesota Pay-As-You-Drive Pricing Experiment

Maya Abou-Zeid*
Massachusetts Institute of Technology, 77 Massachusetts Avenue, Room 1-249, Cambridge, MA 02139
Tel: 617-253-1111, Fax: 617-252-1130, E-mail: mga@mit.edu

Moshe Ben-Akiva
Massachusetts Institute of Technology, 77 Massachusetts Avenue, Room 1-181, Cambridge, MA 02139
Tel: 617-253-5324, Fax: 617-253-0082, E-mail: mba@mit.edu

Kevin Tierney
Cambridge Systematics, Inc., 100 Cambridge Park Drive, Suite 400, Cambridge, MA 02140
Tel: 617-354-0167, Fax: 617-354-1542, E-mail: ktierney@camsys.com

Kenneth R. Buckeye
Minnesota Department of Transportation, 395 John Ireland Blvd., St. Paul, MN 55155
Tel: 651-296-1606, Fax: 651-296-3019, E-mail: kenneth.buckeye@dot.state.mn.us

Jeffrey Buxbaum
Cambridge Systematics, Inc., 100 Cambridge Park Drive, Suite 400, Cambridge, MA 02140
Tel: 617-354-0167, Fax: 617-354-1542, E-mail: jbuxbaum@camsys.com

Total Word Count: 8,695 (Includes 5 Tables, 5 Figures)

* Corresponding Author

Submitted on August 1, 2007 to the Transportation Research Board

Resubmitted on November 16, 2007

ABSTRACT

The Minnesota Department of Transportation (Mn/DOT) carried out a pay-as-you-drive (PAYD) demonstration simulating the replacement of the fixed costs of vehicle ownership and operation with variable costs that give drivers explicit price signals about travel decisions and alternatives. The objective was to estimate the reduction in mileage due to the mileage-based pricing scheme. The study consisted of market assessment surveys and a field experiment. The experiment is the focus of this paper.

The experimental design divided participants into three groups: a control-only group, a treatment-then-control group, and a control-then-treatment group. Participants in the treatment phase were subjected to per-mile prices, and the mileage of all participants was recorded for the entire study duration. Two types of analyses were conducted. Aggregate analyses using bootstrap methods to determine groupwise changes in mileage showed that participants reduced their mileage when charged on a per-mile basis with the greatest reduction during the summer period when trips could be more discretionary in nature. In addition, in order to better understand the variance in mileage sensitivity to per-mile prices, disaggregate analyses were performed using a matching method that matched members of the treatment group to those of the control group based on the probability of participation in the experiment and their baseline mileage. The resulting percentage change in mileage was regressed against percentage change in price and lifestyle variables. The price elasticity of peak period mileage was found to be negative. However, in both aggregate and disaggregate analyses, the price effect was statistically insignificant due to the small sample size.

Key Words: pay-as-you-drive, pricing, traffic congestion, matching

Word Count: 254

1. INTRODUCTION

Pay-as-You-Drive (PAYD) products have emerged as a public policy tool to help in reducing or managing auto travel (*1, 2, 3, 4*). The idea behind PAYD is to transform sunk insurance or ownership costs to costs that vary with automobile usage and for which drivers can receive more frequent price signals (real-time, daily, monthly, or quarterly), resulting in different transportation decision-making, and ultimately less driving. This change in behavior could reduce the need for either costly public sector investments in transportation infrastructure, or invasive public policy programs which seek to alter people's behavior. People would make monthly ownership/leasing and/or insurance payments based on a combination of a fixed monthly charge and a variable rate based on miles driven. In this way, people would receive the same kind of usage-based pricing signals that they today receive from their typical home utility bills. Data collection and transmittal devices now exist which can be installed in vehicles to automatically collect and transmit usage data in the aggregate, or disaggregated by such factors as time of day or route choice. Thus, drivers could be charged not only based on how many total miles they drive, but on whether they drive at peak or off-peak times, or along congested or uncongested routes. In this way, peak-period/peak-route usage could be charged at higher rates, much like utilities do today.

The Minnesota Department of Transportation (Mn/DOT) recently conducted a study to test the feasibility of and interest in PAYD products in the Twin Cities of Minneapolis and St. Paul. The objectives of this study were to (1) simulate the replacement of the fixed costs of vehicle ownership and operation with variable costs that give drivers explicit price signals about travel decisions and alternatives; (2) develop the best possible understanding of transportation price elasticities and how they vary by vehicle ownership/lease arrangement, income, location, annual mileage driven, and other factors; (3) develop an understanding about driver acceptance of use-based fees and appropriate price signals necessary to affect travel behavior changes; and (4) identify strategies and recommendations that might be employed to mainstream or institutionalize policies or techniques learned from the demonstration.

The study consisted of two elements: market assessment surveys and a field experiment. Market assessment and stated preference surveys were conducted to estimate the level of interest in pay-as-you-drive products, the nature of the market for the concept, the response of drivers to price signals (price elasticities) that are based on miles driven, and the overall effect of the program on vehicle-miles traveled and traffic congestion. The findings from the market research have been described previously (*5, 6*).

The purpose of this paper is to describe the field experiment component of this study. The remainder of the paper is organized as follows. Section 2 describes the experiment design. Section 3 describes its implementation. Section 4 presents the data analysis and findings from the experiment. Section 5 concludes the paper.

2. EXPERIMENT DESIGN

The objective of the study is to measure user response to the pricing scheme. Experiments of a similar type where the effect of an intervention is to be evaluated have been commonly designed by grouping participants into a treatment group and a control group, and comparing the behavior of the two groups. An effective approach for reducing seasonal effects and other biases is to divide the participants into two groups (*7*). In an initial experimental period, the first group will be in treatment while the second group will be in control. In a subsequent experimental period, the two groups will be switched; those who were initially in treatment will become in control, and vice versa.

The advantage of this experimental setup, as opposed to having one treatment group and one control group throughout the entire experimental period, is that it clearly isolates the effect of the pricing which should be observed in both experimental periods. Moreover, since every individual is once in treatment and once in control, this allows the identification of the treatment effect within subjects as the individual fixed effects will be controlled.

The experiment design was done based on this “switching” idea and called for the collection of driver data during both control and experiment conditions during an 8-month period, as follows:

- One hundred thirty (130) households were recruited using a random digit dialing technique from households in the eight-county Minneapolis/St. Paul metropolitan area.
- Of these, 30 households were randomly assigned to the “control group.” Their mileage would be tracked over the course of the experiment, but they would not be subjected to pricing experiments. This group was designated as control-control-control (CCC) (the three symbols of the group name refer to three intended periods: a period where all study participants would not be subjected to pricing and two subsequent periods where pricing would take place with switching of the treatment group).
- After all participants would drive for two months while being monitored with electronic devices called CarChips, one-half of the 100-household experiment group (50 households) would be given a pricing experiment. This group was designated as control-experiment-control (CEC). The other half would remain with no pricing. Priced households would drive for three months with simulated prices, and then would go back to not being priced (but being monitored) for the final three months.
- At the beginning of the sixth month, the other 50 experiment-group participants that were still not priced would be given pricing experiments. This group was designated as control-control-experiment (CCE).

Although the experiment was designed to be conducted over an 8-month period, implementation issues arose that required a 3-month extension of the experiment, as described in the next section.

3. EXPERIMENT IMPLEMENTATION

Electronic devices called CarChips that recorded vehicle usage (including mileage by time-of-day and day-of-week) were delivered to participants with installation instructions. Some participants were asked to use CarChips on all household vehicles. Others were asked to install them on only one vehicle. This allowed some participants to substitute mileage on one vehicle to another. Participants were asked to track and report odometer readings for the non-instrumented vehicles in the household so that the impact of this vehicle substitution could be measured. Participants were periodically sent new CarChips and were asked to send in the CarChips that had been installed.

The experiment was conducted by giving each participant household a monetary budget and a rate for each mile driven. Mileage budgets were set based on the number of miles driven during the first month of travel with the CarChip when all vehicles were in a control period. Any money left in the budget at the end of the experiment was theirs to keep. Pricing protocols were assigned randomly and ranged from $0.05 to $0.25 per mile. Pricing for some households was varied for peak and off-peak travel. Ten households were charged a flat fee of $0.05 per mile; 22 households were charged a fee of $0.10 per mile; 11 households were charged a fee of $0.15 per mile; 10 were charged a fee of $0.20 per mile. The remainder of households were charged higher rates in the peak period than in the off-peak periods.

All households remained in control conditions for a second month in order to maintain continuity and to ensure that participants gained experience in swapping car chips. After that, participant households were subjected to different control and pricing regimes. Figure 1 shows the control and pricing period schedules for the different groups of participants. The final schedule covers an 11-month period, and therefore all four seasons. Gasoline prices and other auto operating costs seemed to be stable over the period of the study.

At the beginning of the second experiment period, a personnel issue on the study team made necessary a 3-month period where everybody was in a control period. At the beginning of this period, the participants who were supposed to be subject to pricing were not provided with the necessary information to participate properly. To address this problem, all participants were placed in a control period for three months. The experiment was extended by another 3 months during which the CCE group was subjected to the pricing. CEC participants were also asked to continue in the study through the extension period. A subset of these participants were asked to enter an additional, shorter experiment period within the last few months of the study (designated as CECe), while others continued in a control period (designated as CECc). Those CEC participants who dropped out of the experiment after 8 months (i.e. did not continue through the extended period) were designated as CECx.

At the end of the experiment, an exit survey was administered to evaluate participant attitudes toward the experiment itself, and toward pay-as-you-drive pricing products. Participating households received an incentive of $100 over the course of the study.

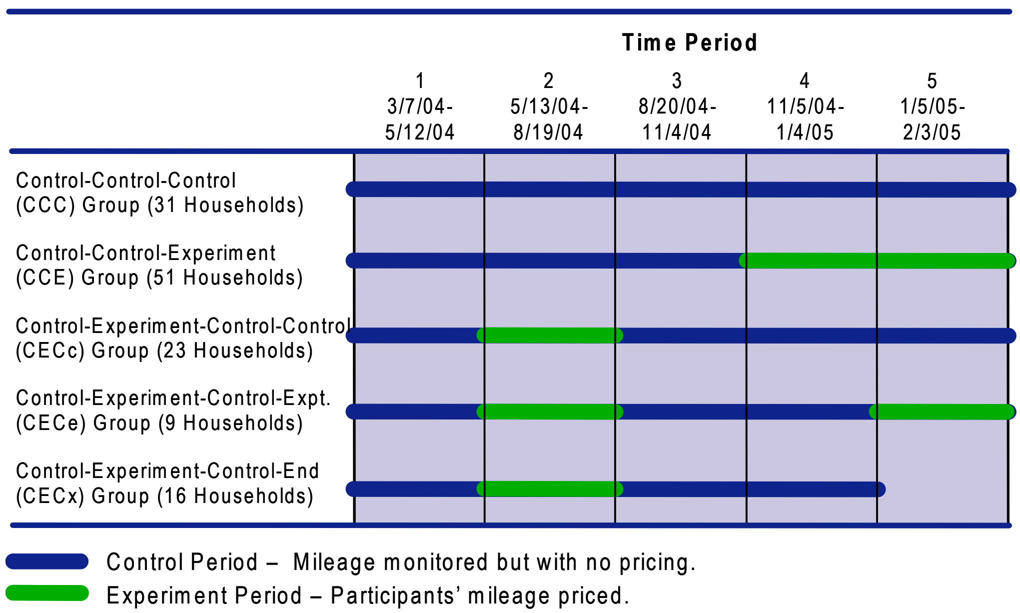

FIGURE 1 Schedule of Control and Experiment Periods by Study Group

4. ANALYSIS

In this section, we present the analysis of the experiment. First, we discuss the respondent characteristics from the recruitment survey. Then we present the evaluation of the pricing effect by discussing the evaluation challenge followed by aggregate and disaggregate analyses. Finally, we describe the conclusions from the exit survey.

4.1 Recruitment Survey

In February 2004, interviewers contacted households in the Twin Cities metropolitan area to collect vehicle usage information and to recruit study participants. There were 2,320 completed surveys for a response rate of 43.1 percent. Most of the 2,320 willing survey respondents were screened out of participating in the experiment for various reasons, such as vehicle availability and CarChip/vehicle compatibility issues. Of those remaining, 660 telephone respondents were asked to participate in the study, and 186 agreed to do so (28 percent). Some later dropped out from the experiment. The demographic characteristics of the cooperating respondents were similar to those who declined to participate in the study, and to those who did not qualify for the study.

Households agreeing to participate had an average number of 2 vehicles and 2 licensed drivers in the household and have lived in the Twin Cities for about 30 years on average. Figures 2, 3 and 4 compare some characteristics of the household participants for the three main respondent groups: Control-Control-Control (CCC); Control-Control-Experiment (CCE); and Control-Experiment-Control (CEC). As these figures show, there were some demographic variations between the groups, but the demographic distributions were similar.

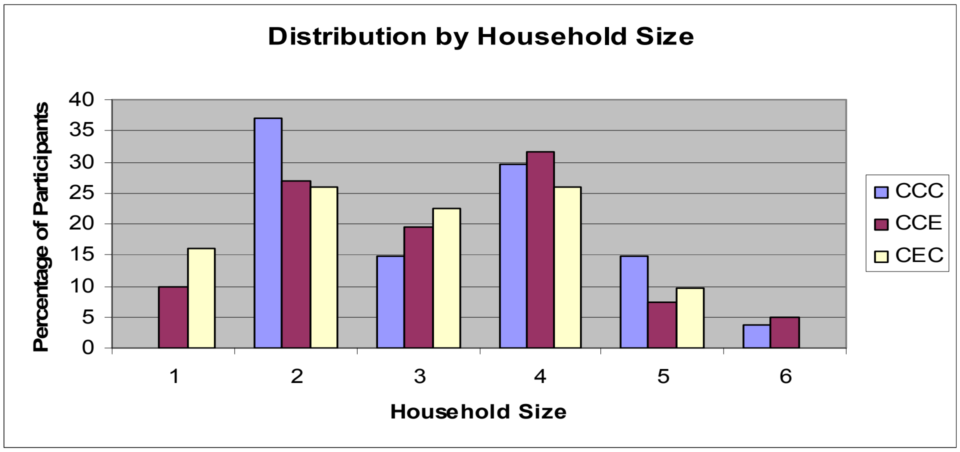

FIGURE 2 Comparison of Household Size Distribution of the Experiment Study Groups
(Sample size: CCC group (N = 27); CCE group (N = 41); CEC group (N = 31))

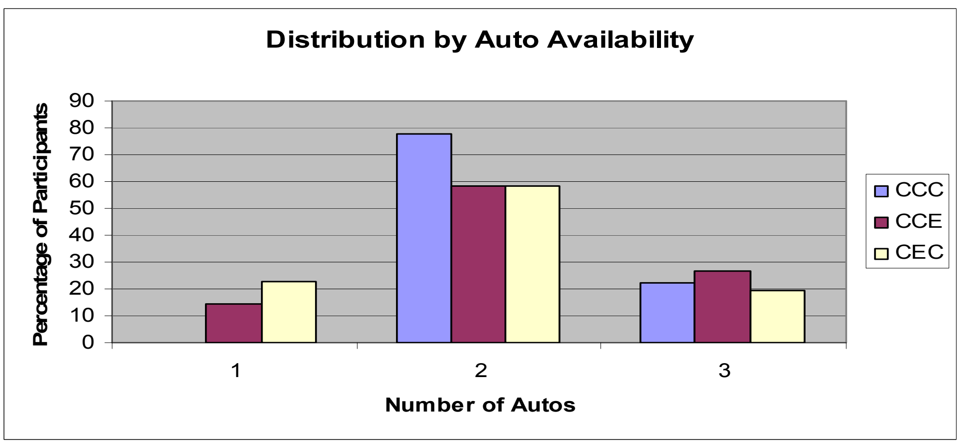

FIGURE 3 Comparison of Number of Household Autos Available of the Experiment Study Groups

(Sample size: CCC group (N = 27); CCE group (N = 41); CEC group (N = 31))

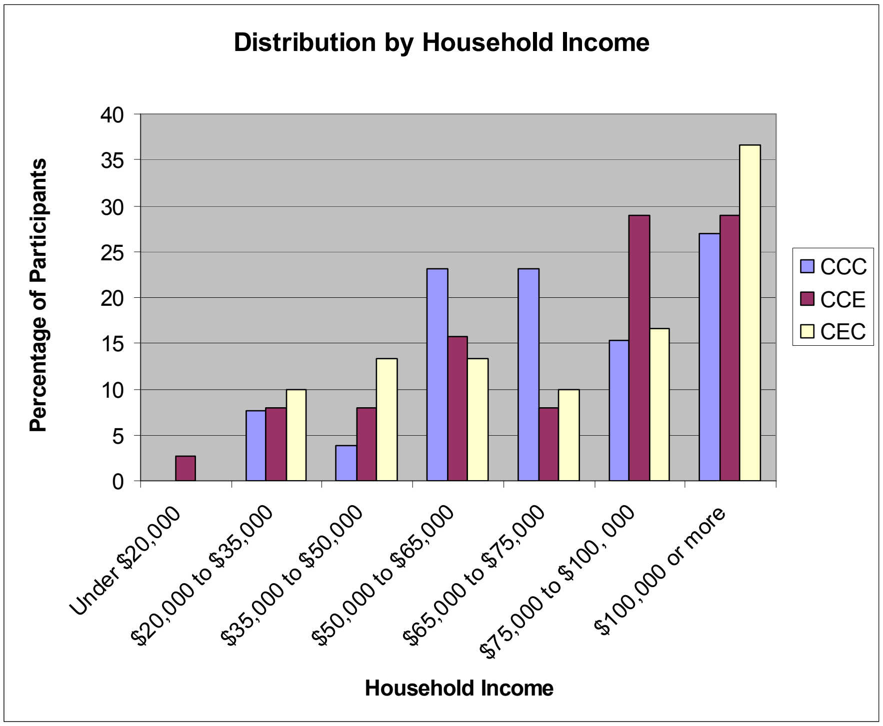

FIGURE 4 Comparison of Household Income Distribution of the Experiment Study Groups

(Sample size: CCC group (N = 26); CCE group (N = 38); CEC group (N = 30))

4.2 Computing the Price Effect

4.2.1 The Evaluation Challenge

A pay-as-you-drive program can be viewed as a social program that aims at reducing mileage through the use of a per-mile price. Evaluation of this program consists of estimating the reduction in mileage due to the price. There is a significant literature on evaluating social programs (*8, 9, 10, 11, 12, 13*). The basic evaluation challenge can be described as follows.

Let Y denote an outcome of interest, and suppose that an individual can be in one of two states: "1" if the individual receives treatment and "0" otherwise. Y_1 is the outcome associated with receipt of treatment, and Y_0 is the outcome in the no-treatment state. The gain of an individual from participating in a program is the change in outcomes between the treatment and no-treatment states, defined as:

$$\Delta = Y_1 - Y_0 . \tag{1}$$

Since at any time an individual can be observed in only one state (treated or untreated), the gain cannot be computed directly for any particular individual. Consequently, the focus in the evaluation literature has been on the estimation of the distribution of impacts among individuals. In voluntary programs and those that target specific groups in the population, the parameter of interest is normally the mean effect of treatment on program participants, defined as:

$$E(\Delta|D=1) = E(Y_1|D=1) - E(Y_0|D=1), \quad (2)$$

where $E(.)$ denotes expected value and D is an indicator of participation ($D=1$ for participants and 0 for non-participants).

The term $E(Y_0|D=1)$ which represents the mean outcome of participants had they not participated, also called a counterfactual, is not observed. The estimation of the desired counterfactual lies at the heart of the evaluation challenge.

4.2.2 Analysis Approaches

The estimation of the treatment effect on the treated can be done using aggregate and disaggregate approaches. Aggregate approaches compare groupwise averages of the variable of interest to determine the treatment effect. Disaggregate approaches estimate the treatment effect for every individual in the treatment sample.

Moreover, estimators can be classified as cross-sectional if the comparison is made between participants and non-participants at one point in time (e.g. in a post-program period) where the non-participant group data are used to estimate the counterfactual; longitudinal if comparisons are made between the same persons in the untreated and treated states (from pre-program and post-program data) where data of participants in the untreated state are used to estimate the counterfactual; and a hybrid of the two if comparisons are made between different persons and using multiple time periods (*14*).

We present below data analyses using both aggregate and disaggregate approaches and cross-sectional and longitudinal estimators.

4.2.3 Aggregate Analysis of Pricing Effect

Figure 5 shows the differences in average daily miles separately for the five distinct experiment time periods. Time Period 1 was unpriced for all participants. The average daily vehicle mileage for this period was 46.7. During Period 2, the average unpriced mileage increased to 49.8 miles, and priced vehicle mileage was reduced to 46.4 miles. In Period 3, there were no pricing data due to the data collection issues, but the average unpriced vehicle had almost the same mileage as the average unpriced vehicle in Period 2. During the fourth and fifth periods, there were almost no differences (statistically insignificant) in the priced and unpriced averages. Compared to the seasonal differences for the unpriced vehicles, the differences between the unpriced and priced vehicles within the same time periods were small.

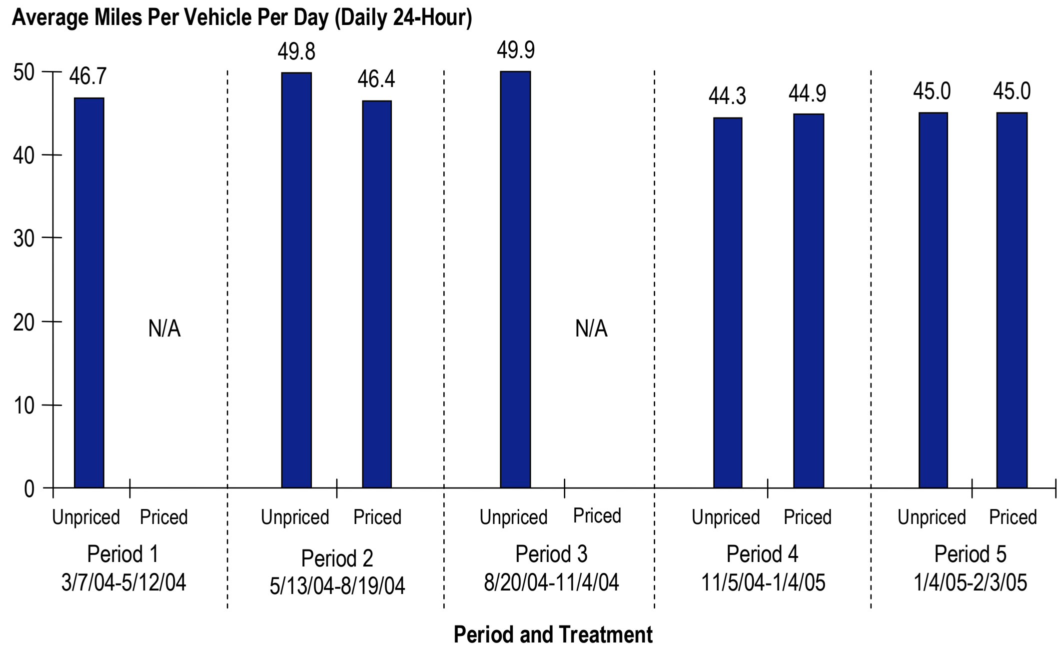

FIGURE 5 Average Miles Per Vehicle Per Day by Calendar Period.

Note: Sample sizes are as follows. Unpriced (N_1 = 109; N_2 = 52; N_3 = 107; N_4 = 62; N_5 = 38); Priced (N_1 = 0; N_2 = 38; N_3 = 0; N_4 = 43; N_5 = 41), where N_t denotes the sample size of a given group in period t.

Another way to evaluate the effect of the pricing treatments is to examine every group of participants separately and evaluate their mileage changes over the different calendar periods of the study. A basic challenge in evaluating these mileage estimates is the variability associated with the small sample size. For this reason, the bootstrap method was used to assess the sensitivity of the estimates to the inclusion or exclusion of households from the analysis. In particular, 50 samples were randomly generated for every group and time period and used to empirically derive average mileage estimates (and differences of these estimates) and standard errors by group and time period.

Table 1 shows the average daily mileage by group and calendar period (without the application of the bootstrap method) with the standard error of the average daily mileage (obtained from the bootstrap method) in parentheses. The mileage pattern of the control-only group (CCC) can be used to track mileage changes that are due to seasonality and possibly other factors caused by unmeasured variables, but not to pricing since this group was not subjected to pricing. This pattern shows that people drive more in the summer (Period 2) compared to the spring (Period 1), and then reduce their mileage again in the fall and winter seasons, with the minimum average mileage occurring between the months of November and January. For the other groups, the mileage changes include both a pricing effect and a seasonality effect possibly combined with other effects due to unmeasured variables.

TABLE 1 Average Daily Mileage by Group and Calendar Period (standard error in parentheses)

			Time Periods		
Group	**3/7/04 to 5/12/04**	**5/13/04 to 8/19/04**	**8/20/04 to 11/4/04**	**11/5/04 to 1/4/05**	**1/5/05 to 2/3/05**
CCC	44.7 (3.2)	51.2 (3.7)	46.0 (2.9)	42.9 (2.4)	44.5 (3.4)
CCE	47.7 (4.3)	48.5 (5.1)	53.5 (3.4)	44.9 (2.3)	46.5 (3.6)
CECc	50.0 (5.0)	47.9 (3.3)	50.8 (4.9)	47.3 (3.4)	45.5 (6.0)
CECe	40.4 (4.4)	42.9 (4.8)	44.7 (5.1)	40.6 (6.1)	38.7 (4.6)

Note: Average mileage is computed based on the following sample sizes: CCC (N_1 = 28; N_2 = 25; N_3 = 28; N_4 = 25; N_5 = 18); CCE (N_1 = 43; N_2 = 27; N_3 = 41; N_4 = 43; N_5 = 33); CCEx (N_1 = 27; N_2 = 27; N_3 = 27; N_4 = 26; N_5 = 20); CCEe (N_1 = 11; N_2 = 11; N_3 = 11; N_4 = 11; N_5 = 8), where N_t denotes the sample size of a given group in period t.

Table 2 shows the difference between each group's average daily mileage in a given time period and the average daily mileage of the CCC group during the same time period with standard error in parentheses. All estimates are obtained using the bootstrap method. Except for the CCE group in period 3, none of the differences between each group's average daily mileage in a given time period when it is unpriced and the average daily mileage of the CCC group during the same time period are statistically significant at the 95 percent level of confidence[1]. Therefore, changes in mileage due to seasonality and other non-price related effects for all groups are assumed to be equal to those of the CCC group.

TABLE 2 Column-wise Difference in Average Daily Mileage: Group Mileage – CCC Mileage (standard error in parentheses) (*)

			Time Periods		
Group	**3/7/04 to 5/12/04**	**5/13/04 to 8/19/04**	**8/20/04 to 11/4/04**	**11/5/04 to 1/4/05**	**1/5/05 to 2/3/05**
CCC	N/A	N/A	N/A	N/A	N/A
CCE	***4.2 (3.6)***	***-2.3 (5.5)***	***8.8 (3.5)***	2.1 (3.5)	3.7 (4.2)
CECc	***4.4 (6.2)***	-3.9 (5.5)	***4.1 (5.6)***	***3.9 (4.0)***	***-0.2 (5.9)***
CECe	***-4.5 (5.8)***	-8.7 (6.2)	***-1.6 (5.7)***	***-3.4 (6.2)***	-5.2 (5.5)

(*) Note: Cell values that are in ***bold and italics*** refer to periods when the corresponding group was not subjected to pricing.

Table 3 shows the difference between average daily mileage for a study group in a given time period (Periods 2 to 5) and the average daily mileage of that group in Period 1 with the standard error of the difference shown in parentheses. Again, all estimates are obtained using the bootstrap method. Except for the CCC group, this difference consists of a seasonality effect (with possibly some other effects due to unmeasured variables) and a price effect. To compute the price effect, we net out the seasonality and other non-price related effects (captured by the

[1] The t-statistics of the difference are computed as the difference divided by its standard error and are compared to a critical t-statistic of 1.96 (normal approximation since the sample size used to compute the standard errors is greater than 50) to assess statistical significance.

change in mileage for the CCC group, since as shown previously, there is no statistical significance between the CCC and other groups) from this difference as shown in Table 4 with standard errors of these differences shown in parentheses. All estimates in Table 4 have also been computed using the bootstrap method.

TABLE 3 Total Row-wise Group Difference in Average Daily Mileage from Period 1: Seasonality and Other Non-Price Related Effects Plus Price Effect (standard error in parentheses) (*)

	Time Periods				
Group	**3/7/04 to 5/12/04**	**5/13/04 to 8/19/04**	**8/20/04 to 11/4/04**	**11/5/04 to 1/4/05**	**1/5/05 to 2/3/05**
CCC	N/A	6.5 (4.6)	1.1 (4.1)	-1.5 (4.2)	-0.3 (4.2)
CCE	N/A	0.0 (6.3)	5.7 (5.2)	***-3.6 (4.5)***	***-0.8 (4.8)***
CECc	N/A	***-1.8 (6.1)***	0.9 (7.6)	-2.0 (5.5)	-4.9 (7.3)
CECe	N/A	***2.2 (7.3)***	4.0 (7.3)	-0.4 (6.4)	***-1.0 (6.1)***

(*) Note: Cell values that are in ***bold and italics*** refer to periods when the corresponding group was subjected to pricing.

TABLE 4 Row-wise Group Difference in Average Daily Mileage from Period 1: Netting out Seasonality and Other Non-Price Related Effects (standard error in parentheses) (*)

	Time Periods				
Group	**3/7/04 to 5/12/04**	**5/13/04 to 8/19/04**	**8/20/04 to 11/4/04**	**11/5/04 to 1/4/05**	**1/5/05 to 2/3/05**
CCC	N/A	N/A	N/A	N/A	N/A
CCE	N/A	-6.5 (6.0)	4.6 (4.7)	***-2.1 (5.2)***	***-0.5 (5.0)***
CECc	N/A	***-8.3 (8.3)***	-0.2 (9.7)	-0.5 (7.5)	-4.6 (6.7)
CECe	N/A	***-4.3 (9.0)***	2.9 (8.3)	1.1 (7.8)	***-0.7 (7.0)***

(*) Note: Cell values that are in ***bold and italics*** refer to periods when the corresponding group was subjected to pricing.

The large standard errors of these differences shown in Table 4 indicate that none of the differences, and consequently the price effects on mileage, are statistically significant at the 95% level of confidence. These large standard errors are due to the small sample size of participants recruited for this study. Budgetary constraints precluded the collection of a larger sample, and future experiments should aim at collecting larger samples to obtain statistically significant estimates.

Despite the statistical insignificance, it should be noted that the pattern of mileage changes due to the pricing makes sense in general. All groups decrease their average mileage during the periods when they are priced. The CCE group reduce their mileage by 2.1 miles per day in Period 4 and 0.5 mile per day in Period 5 (priced periods for CCE). The CECc group reduce their mileage by 8.3 miles during Period 2 which is priced, and the CECe group reduce their mileage by 4.3 miles in Period 2 and 0.7 mile in Period 5 (priced periods for CECe). Based on the group analysis, the CEC groups seem to be more responsive to the pricing treatments than the CCE group. It is likely that the ability to reduce travel is seasonal, with a greater percentage of discretionary trips in the summer. One would assume that these discretionary trips are more likely to be foregone with the pricing incentive in effect. It may also be the case that some of the reduction in driving during the warmer months can be attributed to alternative transportation which might be considered by many to be a more reasonable option during that time of year, since warmer weather and longer daylight generally improve walking, cycling, and transit waiting conditions. However, as mentioned earlier, it is hoped that future experiments will corroborate these results through larger samples.

4.2.4 Disaggregate Analysis of Pricing Effect

We conducted a disaggregate analysis of the pricing effect using the matching method. This method belongs to a broader class of methods which use a comparison group, usually of eligible non-participants, to estimate the outcomes of participants in the no-treatment state (i.e. the counterfactual $E(Y_0|D=1)$). That is, the effect of treatment on the treated is estimated as follows:

$$\hat{E}(\Delta|D=1) = E(Y_1|D=1) - E(Y_0|D=0). \tag{3}$$

The estimator of the treatment effect given by Equation (3) will be an unbiased estimator of the true effect (given by Equation (2)) if the expected value of the outcome Y conditional on D is equal to the unconditional expected value. In other words, the decision to participate should be exogenous in order to obtain an unbiased estimator.

The Matching Method: The method of matching computes the mean effect of a treatment by matching the units (households) in the treatment sample to other nontreated units in a comparison sample and then computing the change in outcomes (mileage) between the matched units. A unit in a treatment group can be matched to one or more units in the comparison (nontreated) group based on similar observed characteristics or on similar probabilities of participation in the program (*15*). The basic assumptions used in matching are that 1) individuals do not enter the program on the basis of gains unobserved by analysts. In other words, it is assumed that the factors that drive participation are observable characteristics of the individual/household, and 2) both treated and nontreated units are available with the same (or similar) observed characteristics X over which the effect of the treatment is to be measured. Given these assumptions, selectivity bias can be removed if one matches units with similar observed characteristics or similar probabilities of participation.

Different matching methods have been developed, such as nearest neighbor matching which assigns one individual with the closest characteristics from the comparison group to match an individual from the treatment group; caliper matching, which matches one individual from the comparison group to one from the treatment group based on a pre-specified tolerance in the difference in characteristics, and; kernel matching, which uses all members of a comparison group with a weighting strategy to match to an individual from the treatment group.

A mean impact estimate based on matching is given by the following expression:

$$m = \frac{1}{N_t} \sum_{i=1}^{N_t} \left(Y_i^t - \overline{Y}_i^c \right) = \frac{1}{N_t} \sum_{i=1}^{N_t} \left(Y_i^t - \sum_{j=1}^{N_c} W(i,j) Y_j^c \right), \tag{4}$$

where the subscripts t and c refer to the treatment group and comparison groups, respectively, N_t and N_c are the sample sizes of the treatment and comparison groups, respectively, Y_i^t and Y_j^c represent outcomes in the treatment and comparison groups, respectively, $\overline{Y}_i^c$ is the weighted average of the comparison group outcome corresponding to individual i of the treatment group, and $W(i,j)$ is the weight assigned to individual j from the comparison group when constructing a match to individual i from the treatment group such that:

$$\sum_{j=1}^{N_c} W(i,j) = 1 \quad \forall i = 1,\ldots,N_t \ . \tag{5}$$

Application to the Pricing Experiment: Applying the matching method to the pricing experiment involves three procedures: 1) developing a participation model; 2) doing the match; and 3) estimating a model of mileage reduction.

Participation Model: Matching the probability of participation reduces the problem of matching to a scalar (one value) instead of matching on a set of observed characteristics. To obtain the probability of participation, the participation model must include all variables that are likely to influence participation.

We used the recruit survey to develop a participation model. The recruit survey includes several household and person variables related to socioeconomics, demographics, and detailed auto characteristics for all households that were eligible to participate even if they chose not to participate. The recruit survey allows us to estimate a model which predicts the probability that a given individual agrees to participate in the experiment. However, several individuals who have agreed to participate later dropped out of the experiment. Therefore, we have developed a sub-model which predicts the probability of not dropping out for any given participant. The two binary logit choice models, "agree to participate" and "not drop out," are then used in combination to compute the probability of participating and not dropping out.

The estimation results indicate that if everything else is the same, respondents are more likely to refuse to participate, but once a household has agreed to participate, the household is more likely to stay in the experiment. Relative to inner counties, households residing in North/East and Southern counties are less likely to participate and more likely to drop out, probably due to the greater reliance on driving in the suburbs as compared to urban areas. As household size increases, respondents are more likely to participate and more likely to stay in the experiment. Fewer autos in the household reduces the likelihood of participation and staying in the experiment probably because it reduces the chance of having alternative unpriced vehicles available to the household. High-income households are less likely to stay in the experiment. The presence of a leased car in the household makes the household less likely to participate and stay in the experiment, while the presence of a shared car among household members increases the likelihood of participation but decreases the likelihood of staying in the experiment.

The effect of annual household mileage on the decision to participate and to stay in the experiment is captured through power series expansions of degrees 6 and 4, respectively. Up to a certain mileage (around 30,000 miles), households are more likely to participate as mileage increases but become less likely to participate as mileage increases further; this is expected because households with high mileage are less likely to benefit from the experiment. Similarly, up to a certain mileage (around 10,000 miles) households are more likely to stay in the experiment as mileage increases but become less likely to stay as mileage increases further because of lower chances for mileage reduction.

Person variables were also included in the models. Females are more likely to participate than males but also more likely to drop out than males. Older people are more likely to participate and stay in the experiment than younger people. Workers are less likely to participate and stay than nonworkers probably because of time constraints. And as people's education level increases, they become more likely to participate and stay in the experiment.

Matching Treatment to Control: The next step after estimating the participation model is to match every household in a treatment group to one or more households in a nontreatment group, such as the experimental control group or a group of eligible nonparticipants. Since mileage data are not available for nonparticipants, we use the experimental control groups as the comparison group from which the matches are drawn. Due to the limited size of the comparison group, we have chosen to use the Kernel matching method so that we can use all observations in the comparison group as a match to a household in treatment.

Since we have three treatment samples corresponding to the Periods 2, 4, and 5, we did the matching separately for each of those three time periods to avoid seasonality effects. For each of these three cases, the comparison group is all households that are in control during that time period.

For every household in a treatment group, we formed a weighted match from the comparison group by assigning a weight to every household of the comparison group so that:

- The weighted probability of participation of the comparison group is equal to the probability of participation of the household in the treatment group;
- The weighted average daily mileage of the comparison group in the first time period is equal to the average daily mileage (in the first time period) of the household in the treatment group; and
- The sum of the weights assigned to all members of the comparison group is 1.0.

The above constraints are applied to estimate the parameters of the weighting function, and the matching is done subsequently.

Results: By matching treatment group members to comparison groups, it was found that while many participants reduced their mileage as expected, several others increased their mileage when subjected to pricing. We postulate, therefore, that the change in mileage is due to the price charged and to variables related to the lifestyle of the household such as indicators of mobility, age, education, presence of kids in the household, etc. The price effect should be negative (i.e. increases in price should decrease the mileage), while lifestyle variables could cause an increase or a decrease in the mileage.

We estimated a linear regression model of the percentage change in peak period mileage as a function of the percentage change in peak period price and of lifestyle variables. The percentage change in peak period mileage is defined as $(M_t - M_c)/M_c$, where M_t is the peak period mileage when in treatment and M_c is the peak period mileage when in control. For a given vehicle in treatment, M_c is obtained from the mileage data of the matched comparison group. The percentage change in peak period price is defined as $(P_t - P_c)/P_c$, where P_t is the cost per mile when in treatment and P_c is the cost per mile when in control. For a given vehicle in treatment, $P_t - P_c$ is equal to the peak period price that the vehicle is charged per mile, and P_c is assumed to be \$0.10 per mile. Although the baseline cost estimate of \$0.10 per mile is arbitrary and could vary over individuals, we did not have information on what the actual costs were for these individuals, and the \$0.10 per mile was deemed as a reasonable average estimate of these costs.

We tried several lifestyle variables retaining only those that were more significant than others and that resulted in a better goodness of fit. The final specification is shown in Table 5. The coefficient of the percentage change in peak period price represents the elasticity of peak period mileage with respect to price. It is negative, as expected, at -0.03 but statistically insignificant as indicated by a t-statistic smaller than 1.96 in absolute value. The only significant lifestyle variables are the base peak and off-peak period mileage in Period 1, which can be thought of as an indication of the mobility of households; as unpriced peak period mileage increases, it becomes more difficult to reduce peak period mileage when priced, hence the positive coefficient of that variable. The coefficient of the off-peak period mileage is negative, signifying a possible substitution effect between the peak and off-peak periods.

TABLE 5 Regression of the Percentage Change in Peak Period Mileage as a Function of the Percentage Change in Peak Period Price and Lifestyle Variables

Variable	Parameter Estimate	t-statistic
Intercept	-0.259	-1.56
Percentage change in peak price	-0.0322	-0.43
Household income (in thousands of dollars)	-0.00106	-0.80
High education dummy	0.0957	1.17
Daily peak period mileage in Period 1	0.0214	5.62
Daily off-peak period mileage in Period 1	-0.00828	-2.42

Note: Number of observations = 94 (few observations with missing income were removed); adjusted R-squared = 0.29

We conclude that the elasticity of peak period mileage is in the right direction, but due to the small sample size, the effect of price on mileage is statistically insignificant, which is the same conclusion reached earlier using the aggregate analysis. Other regressions were tried for both the peak and off-peak change in mileage, but the price effect was again insignificant.

4.3 Exit Survey

At the conclusion of the pricing study, all participants were asked to complete a survey that covered the conduct of the study, their behavior during the study, attitudes toward travel, and their assessments of pay-as-you-drive leasing and insurance concepts.

The survey asked respondents to provide evaluations of different elements of the study. Among the control group there was no basis for change in travel patterns due to pricing during the study. Ninety-three percent of the control group, versus 69 percent of the experiment group, agreed that the study did not affect their driving habits.

Both experiment and control groups felt that price uncertainty would be an important factor in considering whether to try pay-as-you-drive insurance and leasing. Another important factor was the potential cost savings. The control group felt that the ability to control costs by reducing mileage was not as important as the experiment group. Compared to the experiment group, the control group felt that privacy concerns were a more important consideration in their adoption of pay-as-you-drive insurance and leasing.

Exposure to the experiment made respondents more receptive to consider alternate modes of insurance and vehicle purchases. Consistently, the experiment group was more likely to choose pay-as-you-drive insurance and leasing if available. In addition, the experiment group was more likely than the control group to consider pay-as-you-drive insurance and leasing if features such as variable mileage pricing by time of day and yearly audits were offered. An overwhelming majority of the participants said that they were more likely to choose pay-as-you-drive insurance if they could switch back to traditional insurance without penalties.

5. CONCLUSION

We presented the findings from an experiment conducted by the Minnesota Department of Transportation to evaluate user response to pay-as-you-drive products. The experiment was part of a larger project simulating the replacement of the fixed costs of vehicle ownership and operation with variable costs that give drivers explicit price signals about travel decisions and alternatives.

The experiment was conducted over a period of approximately one year. Participants were divided into three groups: a control-only group, a treatment-then-control group, and a control-then-treatment group. Participants were then given a budget at the beginning of the experiment and charged a per-mile price during the treatment phase of the experiment. Participants could save money by driving less. Mileage was monitored through electronic devices called CarChips as well as through odometer readings.

The experimental data were analyzed both at the aggregate level using groupwise comparisons and bootstrap analysis methods and at the disaggregate level by matching treatment group members to similar members in a comparison group to compute the pricing effect. While it was found that some households increased their mileage when in treatment, both types of analyses reflected an overall average reduction in mileage when the mileage-based prices are used, as also confirmed during the exit survey. However, the analysis of the data was limited by the small sample size and by the problems encountered during the data collection phase. The small sample size caused the effect of price on mileage to be statistically insignificant. Future experiments would benefit from recruiting a larger sample of participants. Yet, the price effect was in the right direction, and it is the conclusion of the study that wide-scale per-mile pricing would result in a measurable, but small, reduction in vehicle mileage. In magnitude, this reduction is probably within the regular variation that occurs from season to season, but if these reductions were generalized to the entire population per-mile pricing would reduce VMT and congestion measurably.

Other analyses, not reported here for space limitations, revealed that on a percentage basis, the biggest reduction in mileage would be on weekends, which presumably have the highest percentage of discretionary travel purposes, but weekday peak-period travel would be reduced by more than weekday off-peak period mileage. Over the course of the study, the average daily mileage of unpriced vehicles was 47.5 miles, compared to an average of 45.4 miles for the priced vehicles (4.4 percent difference). Comparatively larger differences in percentage terms were measured for weekend trips (8.1 percent) and for weekday peak-period trips (6.6 percent). In addition, mileage reductions from per-mile pricing would vary by season, with the highest reductions during the warmer months. The reader is referred to (*16, 17*) for the additional analyses and policy implications of this study.

ACKNOWLEDGEMENTS

The authors would like to acknowledge John Schamber and David Bender of MarketLine Research in Minneapolis for fielding the market research instruments. Jean Wolfe, Billy Bachman and others at Geostats handled the deployment of the technology, management of the participants and retrieval and quality review of the experiment data. Kimon Proussaloglou contributed to the research design. Karen Hamilton assisted in preparation of the synthesis of the market research data and reporting. Krishnan Viswanathan and Casey Frost contributed to project implementation, and Marc Cutler provided oversight to the consultant work. Funding for the project came from FHWA's Value Pricing Pilot Program and Mn/DOT.

REFERENCES

1. "Project XL: Progressive Auto Insurance." U.S. Environmental Protection Agency. Last updated on March 2, 2006. Available online: http://www.epa.gov/projectxl/progressive/

2. Vickrey, W. Automobile Accidents, Tort Law, Externalities and Insurance: An Economist's Critique. 33 *Law and Contemporary Problems*, pp. 464-470, 1968.

3. Funderberg, K., M. Grant, and E. Coe. Changing Insurance One Mile at a Time. *Contingencies* (www.contingencies.org/novdec03/changing.pdf), Nov./Dec. 2003, pp. 34-38.

4. Litman, T. Distance-Based Vehicle Insurance as a TDM Strategy. *Transportation Quarterly*, Vol. 51, No. 3, Summer 1997, pp. 119-138. An updated version of this paper is available at VTPI.

5. Buckeye, K.R., J. Buxbaum, M. Abou-Zeid, and K. Tierney. Minnesota Pay-As-You-Drive Market Research. Presented at the 86th Annual Meeting of the Transportation Research Board, Washington, D.C., 2007.

6. Cambridge Systematics, Inc. *Mileage-Based User Fee Demonstration Project: Market Assessment Survey Results*. Final report 2006-39B, Minnesota Department of Transportation, 2006. Available online: http://www.lrrb.org/200639B.pdf

7. Fehr, E., and L. Göette. Do Workers Work More if Wages Are High? Evidence from a Randomized Field Experiment. Working paper No. 125, Institute for Empirical Research in Economics, University of Zurich, 2005.

8. Heckman, J. J., H. Ichimura, and P. Todd. Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Program. *Review of Economic Studies*, Vol. 64, 1998, pp. 605-654.

9. Heckman, J., and R. Robb. Alternative Methods for Evaluating the Impact of Interventions. In *Longitudinal Analysis of Labor Market Data*, ed. James Heckman and Burton Singer, 1985, pp. 156-245. Cambridge: Cambridge University Press.

10. Heckman, J. J., and J. A. Smith. Assessing the Case for Social Experiments. *Journal of Economic Perspectives*, Vol. 9, No. 2, 1995, pp. 85-110.

11. Heckman, J. J., and J. A. Smith. The Pre-Program Earnings Dip and the Determinants of Participation in a Social Program: Implications for Simple Program Evaluation Strategies. NBER Working Paper No. W6983, 1999.

12. Parker, S. W., and G. M. Teruel. Randomization and Social Program Evaluation: the Case of Progresa. 2003. Available online:

http://www.campbellcollaboration.org/Bellagio/Papers/CRT_Parker2.pdf

13. Zhu, S., C. M. Anderson, G. J. Tedeschi, B. Rosbrook, C. Johnson, M. Byrd, and E. Guitiérrez-Terrell. Evidence of Real-World Effectiveness of a Telephone Quitline for Smokers. *The New England Journal of Medicine,* Vol. 347, No. 14, 2002, pp. 1087-1093.

14. Heckman, J. J., R. J. LaLonde, and J. A. Smith. The Economics and Econometrics of Active Labor Market Programs. *Handbook of Labor Economics*, Vol. 3, Ashenfelter, A. and D. Card, eds., Amsterdam: Elsevier Science, 1999.

15. Rosenbaum, P., and D. Rubin. The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika*, Vol. 70, No. 1, 1983, pp. 41-55.

16. Cambridge Systematics, Inc. *Mileage-Based User Fee Demonstration Project: Pay-As-You-Drive Experimental Findings*. Final report 2006-39A, Minnesota Department of Transportation, 2006. Available online: http://www.lrrb.org/200639A.pdf

17. Cambridge Systematics, Inc. *Mileage-Based User Fee Demonstration Project: Potential Public Policy Implications of Pay-As-You-Drive Leasing and Insurance Products*. Final report 2006-39C, Minnesota Department of Transportation, 2006. Available online: http://www.lrrb.org/200639C.pdf