

Is Econometrics Useful for *Private* Policy Making?

A Case Study of Replacement Policy at an Auto Rental Company

Sungjin Cho, *Seoul National University*
John Rust, *University of Maryland**

April 22, 2008

Abstract: The goal of this paper is to illustrate the *potential* usefulness of econometrics as a tool to assist *private* policy makers. We provide a case study and detailed econometric analysis of the automobile replacement policy adopted by a large car rental company. Unlike public policy making — where the benefits from using econometric models and “science-based” approaches to policy making are hard to quantify because the outcomes of interest are typically subjective quantities such as “social welfare” — in the case of firms there is an objective, easily quantifiable criterion for judging whether policy A is better than policy B: *profits*. We introduce and estimate an econometric model of the rental histories of individual cars in the company’s fleet. Via stochastic simulations, we show that the model provides a good approximation to the company’s actual operations. In particular, the econometric model is able to reproduce the extraordinarily high rates of return that the company obtains on its rental cars, with average internal rates of return between purchase and sale of approximately 50%. However, the econometric model can simulate outcomes under a range of counterfactual vehicle replacement policies. We use the econometric model to simulate the profitability of an alternative replacement policy under pessimistic assumptions about the rate maintenance costs would increase and rental rates would have to be decreased if the company were to keep its rental cars longer than it does under the *status quo*. Depending on the vehicle type, we find that the company’s expected discounted profits would be between 6% to over 140% higher under the suggested alternative operating strategy where vehicles are kept longer and rental rates of older vehicles are discounted to induce customers to rent them. The company found this analysis to be sufficiently convincing that it undertook an experiment to verify the predictions of the econometric model.

Keywords: automobile rental/leasing, duration models, semi Markov processes, optimal replacement policy, optimal stopping problem, dynamic programming, policy evaluation

JEL classification: C13-15

*Corresponding author: John Rust, Department of Economics, University of Maryland, 3105 Tydings Hall, College Park, MD 20742, phone: (301) 404-3489, e-mail: jrust@gemini.econ.umd.edu. Sungjin Cho can be reached at Department of Economics, Seoul National University, Seoul, Korea, phone: 82-2-880-6371, email: sungcho@snu.ac.kr. We are grateful to Adam Copeland of the Bureau of Economic Analysis, the editors, and a referee and research assistants Cyrus Aghamolla and Craig Drohan for helpful comments on a previous version of this paper that lead to significant improvements. We are especially grateful to the rental car company we studied, and the person in this company who provided us with data and took the time to evaluate our findings. Unfortunately, confidentiality restrictions prevent us from revealing this person or the name of the corporation in question.

1 Introduction

The goal of this paper is to illustrate the *potential* usefulness of econometrics as a tool to improve policies chosen by *private* policy makers. We provide a case study and detailed econometric analysis of the vehicle replacement policy of a large rental car company. Contrast this with public policy making where the benefits from using econometric models and “science-based” approaches to decision making are hard to quantify because the outcomes of interest are typically subjective quantities such as “social welfare.” However in the case of firms there is an objective, easily quantifiable criterion for judging whether policy A is better than policy B: *profits*. Thus, one way to demonstrate the value of econometric models is to use them to help firms increase their profits.

In a related article, Rust (2007) presents a separate case study that focuses on whether econometrics is useful for improving *public* policymaking. That paper draws on Rust’s experience as a former consultant to the U.S. Social Security Administration. Unfortunately in that application, which involved predicting the welfare and cost/benefit impact of a proposed change in the U.S. Disability Insurance system, Rust concludes that the prospect that econometric models will have any measurable role or impact on public policymaking is bleak — at least in the short term under the current regime in Washington, DC where “faith based” policymaking and political considerations typically dominate scientific advice. In our view, the best chance to interest public policymakers in the value of econometric models in the long run is to start by providing clear-cut demonstrations of the usefulness of econometric models in improving policymaking in private sector applications in the short run. While the issue of rental policy is admittedly less interesting, ambitious, or important *per se* than Social Security policy, it is also far less politically sensitive. Also, because firms often have stronger incentives to improve outcomes and have more control over their own actions than do government bureaucrats, there is a greater chance that private policymakers will pay attention and actually take concrete actions in response to predictions and policy recommendations from econometric models than generally is the case in public sector applications.

We illustrate these points using a new data set on rental car histories that we obtained from a major rental car company. These data provide a unique opportunity to test the hypothesis that this company has adopted a profit maximizing vehicle replacement policy. We are not aware of previous studies that have questioned the widely adopted policy by rental car companies of replacing their cars frequently. Indeed many major rental car companies (e.g. Hertz, National, etc.) typically sell cars after only 20,000 miles. However, this is an extremely costly policy due to the well known rapid depreciation in car prices.

The econometric model predicts that the company can significantly increase its profits by keeping its rental cars longer than it currently does – even under pessimistic or “worst case” assumptions to be described below. We build on results from Cho and Rust (2007) who used numerical dynamic programming to calculate optimal replacement policies under these worst case assumptions. Their findings indicate that profits can be significantly increased if vehicles are replaced after roughly 90,000 miles (or about 5 years) instead of 3 years (or about 45,000 miles) under the company’s current replacement policy. We use the econometric model to conduct stochastic simulations of the company’s profits under this suggested counterfactual replacement policy. Depending on the vehicle type, the simulations predict that expected discounted profits would be 6% to over 140% higher, depending on the type of vehicle. These predictions were made under pessimistic assumptions about the degree maintenance costs will increase and the amount rental rates of older vehicles would have to be decreased (to induce customers to rent them) if the company were to keep its vehicles significantly longer than it currently does. The company found this analysis to be sufficiently convincing that it decided to undertake an experiment to verify our prediction that keeping its cars longer would significantly increase its profits.

This study also illustrates the limitations of simple “reduced-form” econometric approaches to testing the profitability of the company’s current rental car replacement policy, similar to the “freakonomics” approach adopted by Levitt (2006). The main value added from constructing a more explicit econometric model of the rental histories of individual cars in the company’s fleet is that a wide range of outcomes can be *simulated* under both the *status quo* and a range of counterfactual replacement policies. This is not possible under Levitt’s approach. Our simulations demonstrate that the econometric model provides good approximation to the company’s actual operations — at least under the *status quo*. In particular, simulations of the econometric model are able to reproduce the extraordinarily high rates of return that the company obtains on its rental cars, with average internal rates of return between purchase and sale of approximately 50% under the company’s current replacement policy.

The methods used in this paper are an extension and specialization of work on optimal replacement of durable goods (Rust 1985, 1987 and Cho 2007). The main difference between the approach taken in this paper and this previous work is that these previous studies *assumed* that firms are behaving optimally, whereas in this study we relax and empirically test this assumption. It is possible for us to relax and test the maintained assumption of profit maximizing behavior in this case precisely because the firm’s objective, *expected profit maximization*, does not involve subjective quantities that need to be estimated. This stands in stark contrast to most political or public policy applications, where the behavior of individuals

or organizations depends critically on what their objectives are. Economists typically presume that these agents are acting “as if” they were maximizing the expected value of well-defined utility function or a social welfare function. However, unlike profits, utility and social welfare functions are typically not known explicitly, even by the decision makers in question. As a result, strong assumptions have to be imposed in order to estimate them, including the hypothesis of *expected utility maximization*. Rust (1994) has shown that the hypothesis of expected utility maximization *per se* is not sufficient to uniquely identify an agent’s underlying utility function. Indeed, there is a fundamental identification problem: it is often possible to “rationalize” many types of behavior as being optimal for some appropriately defined objective function. The identification problem is another reason why it is so difficult to apply econometric models to help improve public policymaking, because it is very hard to be sure what the “true objectives” of citizens and policymakers really are, and given incomplete knowledge of payoffs/rewards and how to aggregate and resolve conflicting objectives of the many different individuals who are affected by public policy decisions, it is hard to provide an objective basis for judging whether policy A is “better” than policy B.

Rental car companies are an ideal testing ground for an econometric approach to policymaking since these companies have large fleets, make frequent replacement decisions, and have good records of their operations and decisions. Furthermore, there is little disagreement that their basic objective is to maximize expected discounted profits, and it is straightforward to measure profits in this application. We feel extremely fortunate to have been able to earn the trust of a large rental car company, which has provided their operational data to us, and allowed us to study and dialog with them. In order to protect the confidentiality of this firm and continue our relationship with them, we cannot provide any further information about the firm or its operations or exchange the data we obtained from the company with other researchers. However we would like to make clear that we have no financial relationship with this company and have not received any monetary compensation for our services. In effect, we are *consulting for data* and the fact that we are doing this work with a scientific rather than financial objective in mind enables us to conduct a more independent and unbiased analysis of the company’s operations.

Section 2 describes the rental car data and Section 3 introduces and estimates an econometric model of the rental company’s operations. In Section 4 we show, via stochastic simulations, that this model provides a good approximation to actual outcomes for this company. Section 5 shows how the econometric model can be used to evaluate counterfactual operating strategies. We show that an alternative replacement policy of keeping rental cars longer results in significantly higher profits. We make some concluding remarks in Section 6.

2 Data and Preliminary Regression Analysis

We obtained data from a large syndicated rental car company that owns and rents a large fleet of rental vehicles. The company provided us with data on over 3900 individual vehicles at various rental locations. These do not represent the entire fleet at any point in time, but they do represent a significant share of the company's holdings. All of these vehicles were first acquired (i.e. registered) after 1999, and almost all of these vehicles were purchased brand new from auto manufacturers. While there are occasional "group purchases" of particular brands and models of vehicles on the same date, when these group purchases did occur, they typically amounted to only 4 or 5 vehicles of the same brand/model at the same time. Thus, this company by in large follows an *individual* vehicle replacement and acquisition strategy, as opposed to "block acquisitions and replacements" i.e. simultaneously acquiring and disposing of large groups of vehicles of the same make and model at the same time.

The data include the date and purchase price for each vehicle the company acquired, the date and odometer value when the vehicle was sold, and the complete history of maintenance and rentals between the purchase and sale dates. The rental contract data record the dates each contract started and ended, and (sometimes) the odometer value of the vehicle at the start and end of the rental contract. We found (with the exception of the odometer value at the date each vehicle was sold, which was accurately recorded), the company's data on odometer values at the beginning and end of each rental contract to be frequently missing or based on guesses by the company's rental agents. This was especially true for long term contracts that were "rolled over" (i.e. where the contract was renewed by the customer without returning it back to the lot). Many of the company's rental agents appear to have filled in rough estimates of the out and in odometer values at the roll over dates when customers informed them of their to keep their car another month. As a result, we did not trust most of the in or out odometer readings in the company's rental records. In order to infer the driving patterns and number of kilometers typically traveled during each rental contract we relied on some (we believe reasonable) econometric modeling assumptions that we will describe shortly.¹

¹ The data include records on the date of accidents and the cost of repairing accident damage, as well as decisions to scrap (versus sell) vehicles that were sufficiently badly damaged as a result of accidents. Although 2543 of the 3908 vehicles in the data set experienced one or more accidents over their service lives, only 123 vehicles were sufficiently badly damaged that they had to be scrapped. In almost all cases where accidents have occurred, the cost of repairing the damage to the vehicle is covered by the insurance of the renter (if the renter was at fault), the insurance of the other party to the accident (if they were at fault), or by the company's insurance (if the party at fault has no insurance). There is a potential indirect source of financial loss due to accidents that the company is not compensated for, namely, if the resale price for cars with accidents is lower. However this effect can be expected to be small, since whenever an accident is repairable, the insurance pays all necessary repairs to restore the car to its pre-accident condition. The company is required to report the number of accidents

The company rents its cars on two types of contracts: a long term contract or a short term contract. Long term contracts are typically written with a maximum duration of one month, combined with a right to automatically renew the previous contract for another month. Rental contracts are at a daily rate with no additional charges for distance traveled during the contract. The daily rate for a long term contract is typically lower than the daily rate for short term contracts. There is a penalty for early returns of vehicles in long term contracts, generally equal to 20% of the lost rental revenue for the unfinished remaining days in the contract.

Figure 1 illustrates typical rental histories for three different cars in the company’s fleet: 1) a compact car rented from one of the company’s urban locations, 2) a luxury car rented from an urban location, and 3) recreational vehicle rented from a “tourist” location. In the graphs, a value of ‘0’ denotes a car that is on the lot waiting to be rented, a value of 1 denotes a long term contract, and a value of 2 denotes a short term contract.

We see that the compact and luxury cars that were rented from the urban location started out in a series of long term rentals, with no intervening “lot spells” between the successive monthly rental contracts. It is possible that a succession of unbroken long term contracts might represent a customer who rolled over their monthly contracts into the *de facto* equivalent of a lease, lasting nearly one year in the case of the compact and two years in the case of the luxury vehicle. After these long term contracts came to an end, these vehicles were rented on a series of short term contracts, except that the compact car was rented for a final long term contract episode for 30 days near the end of its service life. On the other hand, all of the rentals of the recreational vehicle (RV) in the tourist location were short term rentals, with most contracts lasting only a few days.

Figure 1 also shows the exact service life and the realized internal rate of return (IRR) that the company earned on the vehicle over its service life. The IRR is defined at the discount rate r (where r is measured on an annual basis) that sets the net present value of the cash flow stream earned by the company over the vehicle’s service life equal to zero:

$$0 = \sum_{t=0}^T \exp\{-a_t r/365\} c_t, \quad (1)$$

where T is the number of days over which cash inflows or outflows occurred for the vehicle, c_t is the cash inflow (if positive) or outflow (if negative), and a_t is the number of days after the initial purchase of the

and information on the nature of each accident (severity, cost of repair and so forth) that a vehicle experienced at the time it is sold. However the econometric evidence we offer below shows that neither the total number of accidents, nor the total cost of repairing these accidents is a significant predictor of resale prices.

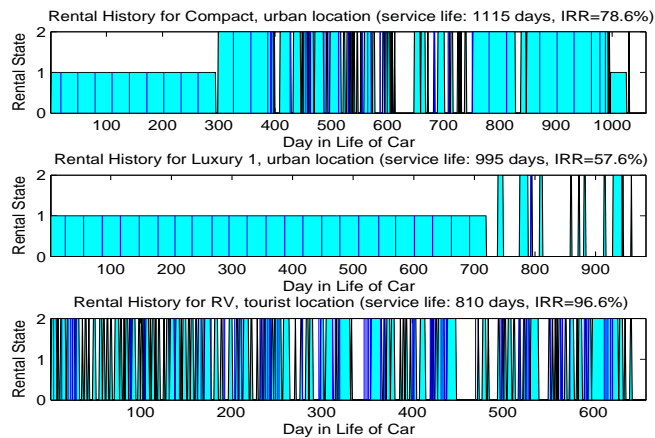


Figure 1 Typical Rental Histories for Three Cars in the Company Fleet

vehicle that the t^{th} cash flow occurred. Thus, $c_0 < 0$ and $a_0 = 0$ represent the initial purchase of the car, and then subsequent cash flows would be rental revenues received when the car returned at the end of each rental contract, and cash outflows for maintenance on the dates they occurred. The final cash flow, $c_T > 0$, is the resale price the company receives from selling the car in the used car market, or at an auction. Thus, a_T represents the service life, i.e. the actual age of the car in days at which it was sold, assuming its initial age was $a_0 = 0$ (since all cars were purchased brand new).

We see that for each of the cars illustrated in figure 1, the realized rates of return are extraordinarily high. The firm earned a 78.6% rate of return on the compact car, a 57.6% rate of return on the luxury car, and a 96.6% rate of return on the recreational vehicle. The undiscounted profits are also high — \$16,683 for the compact, \$24,753 for the luxury car, and \$27,654 for the RV — especially in relation to the initial purchase prices of these cars: \$9011, \$22808, and \$17889, respectively. The odometer values (in kilometers) on these cars at time of sale were approximately the same, 66300, 61000, and 63265, respectively. However the depreciation rates experienced in the resale values (i.e. the ratio of resale price to new price) of the three cars was quite different: 39%, 56%, and 56%, respectively. The fact that the compact experienced relatively greater price depreciation could be due to having been driven longer (with a terminal odometer of 66300 and service life of 1115 days, it was approximately 10% older at time of sale than the other two vehicles), or it could just reflect a lower level of durability and thus a greater level of price depreciation for any given odometer value.

The internal rates of return earned by these three example cars in figure 1 are not atypical: the mean IRRs for all compacts, luxury, and RV's of the same make, model and vintage as these were 77%, 49%,

and 53%, respectively. Table 0 presents the results of a regression of the internal rate of return on various explanatory variables to see which factors are most important predictors of high returns for rental vehicles. We report three regressions for the vehicle types: compact, luxury, and recreational vehicle, pooling over all rental locations.² The predicted signs of the coefficients are mostly consistent with intuition: the utilization rate should have a positive coefficient for reasons discussed above, maintenance costs and the new purchase price should have negative coefficients, the sale price should have a positive coefficient, and the daily rental rates for long and short term rental rates should have positive coefficients.

Variable	Compact All Locations		Luxury All Locations		RV All Locations	
	Estimate	(t-stat)	Estimate	(t-stat)	Estimate	(t-stat)
Constant	0.575	(2.33)	-0.006	(-0.02)	0.999	(1.74)
Utilization Rate	0.003	(0.02)	0.522	(4.14)	1.366	(4.52)
Fraction Rented Long Term	-0.220	(-2.52)	-0.076	(-0.88)	-0.876	(-4.79)
Total Maintenance costs (\$000)	$-7.46e^{-5}$	(-2.81)	$-2.00e^{-5}$	(-1.64)	$6.978e^{-6}$	(0.22)
Odometer (000 km)	0.0007	(1.23)	-0.0004	(-0.56)	-0.001	(-0.79)
Age at Sale (years)	0.151	(3.98)	0.072	(1.54)	-0.154	(-1.11)
New Price (\$000)	-0.104	(-4.98)	-0.036	(-4.22)	-0.082	(-3.22)
Sale Price (\$000)	0.008	(0.37)	-0.002	(-0.23)	0.063	(1.68)
Short term rental rate	0.003	(2.83)	0.0006	(2.07)	0.004	(3.84)
Long term rental rate	0.037	(19.09)	0.020	(7.25)	0.009	(1.12)
Observations, R^2	167	0.806	40	0.776	31	0.859

Table 0: Regression Results for Dependent Variable IRR

However we are unable to draw any clear conclusions about the effect of age and odometer value on the IRR: in some cases the coefficients of these variables are positive, and in others negative, and the coefficient estimates are generally statistically insignificant. The results for the maintenance cost variable are also ambiguous. There are a number of possible reasons why the coefficients of age, odometer, and maintenance costs have variable signs and are frequently statistically insignificant. One reason is that these variables have a high degree of collinearity, especially age and odometer. When we re-run the regressions and include only age or odometer individually, the results are still ambiguous, and the coefficients are generally statistically insignificant. Only in one case, for the luxury vehicle, are both age and maintenance statistically significant when odometer is omitted from the regression, and we see in Table 0 that age has a positive coefficient and total maintenance cost has a negative coefficient. But even in this case, the effect

²The results are basically unchanged if we use the logarithm of the internal rate of return as the dependent variable: the R^2 statistics are slightly lower but the same pattern of signs and significance levels for the coefficients emerges for this alternative specification for the dependent variable in the regression.

of age on IRR is small: the regression results predict that keeping a luxury car for 100 more days increases the IRR by only 0.03 (i.e. by about 5% of the mean IRR of 0.533).

One potential interpretation of the small and statistically insignificant coefficients on age and odometer is that it is an indication of optimizing behavior by the firm. That is, if the firm is choosing age and/or odometer value approximately optimally, we would expect that any variations in these variables about their optimal values should be small. Let $\Pi(o)$ denote the expected discounted profits from keeping a car until it reaches the odometer threshold o before selling it. The optimal odometer threshold o^* is the solution to

$$\frac{\partial \Pi}{\partial o}(o^*) = 0. \quad (2)$$

It follows that if the odometer values at which the company sells its cars are approximately equal to the optimal threshold o^* , we would not detect any significant effect on discounted profits from small variations in the realized odometer value about its optimal value o^* at the time the car is sold. Since IRR is monotonically related to discounted profits, it follows that if the firm is behaving approximately optimally, the effect of small deviations in odometer value from o^* on IRR should also be approximately zero.

However there are a number of reasons why this interpretation may not constitute convincing evidence of optimal behavior on the part of the company. First, as we will show in the next section, the range of odometer values at which the company replaces its vehicles is very wide, more than 100,000 kilometers wide. The argument we made above will only be valid for relatively small deviations of o from its optimal value o^* , and a 50,000 kilometer deviation on either side of o^* seems too large for our argument to apply.³

Indeed, as we will show in section 5, the optimal threshold is generally not a single value o^* , but rather a *function* $o^*(d, r)$ that depends on the rental state of the car r and the duration d in this state. Under the optimal policy (which we calculate by numerical dynamic programming) there is a wide range of odometer thresholds at which replacement can be optimal depending on the values of (d, r) . While it is true that small deviations of replacements about $o^*(d, r)$ have small effects on *infinite horizon discounted profits*, it does not follow that such deviations necessarily have small effects on the internal rate of return of the *current car*, since this ignores the effects of delaying replacing the current car on the stream of profits from the car that replaces the current car, the one after that, and so on into the infinite future. Clearly, if we fail to consider that the firm is interested in maximizing expected discounted value of the *stream of current and future profits* and not just the profits on the currently held vehicle, we can get very misleading results.

³The argument could also be made that the company is choosing an optimal replacement age a^* , but as we will see, there is also a wide range of ages over which the company replaces its vehicles. So the same problem would apply if we hypothesized that the company's replacement threshold was defined in terms of vehicle age rather than odometer value.

In particular, keeping the current car longer will nearly always increase profits on the margin since even for older cars incremental maintenance costs from keeping a current car a little bit longer are generally far lower than the incremental rental revenues, but keeping the current car longer comes at an opportunity cost in terms of higher profits that might be earned by replacing a currently held old car with a new one. We need to calculate the infinite horizon profit function (i.e. the value function) to properly account for this trade-off, and this is why a naive approach to calculation of an optimal replacement threshold o^* in equation (2) above is likely to be misleading.

There is also reason to believe that the coefficient estimates for age and odometer in Table 0 are untrustworthy because these variables are *endogenous*. That is, the company's replacement decisions clearly determine how old and how high the odometer is on its vehicles before they are replaced. If there are unobserved factors associated with a car that lead it to be more profitable (i.e. have higher IRR) these same factors could also lead the company to want to keep the car longer. As a result, one might expect that age and odometer to be positively correlated with unobserved factors affecting profitability and IRR, and this correlation can lead to a spurious upward bias in the coefficient estimates for age and odometer value.

As a result, it is difficult to draw any firm conclusions from Table 0 about whether the company is behaving approximately optimally or not. We would need some sort of *instrumental variable* to deal with this endogeneity problem, but there are no obvious candidates for valid instruments in the data set. What we want would be one or more variables that resulted in *exogenous* shifts in the age at which the company replaced some of its vehicles. An example of such a variable might be a *recall variable*, that is, if there was some major problem in one of the types of cars that the company owned that leads to a recall to the manufacturer, or convinces the company to sell these vehicles before it had intended to sell them. In such case, the "premature" sales of the vehicles could be regarded as a "quasi experiment" that could provide information on how exogenous reductions in vehicle age or odometer values at time of sale would affect the IRR. Unfortunately, there were no recalls of vehicles in the dataset.

The only alternative instrument that we are aware of is an *accident dummy*. If we assume that accidents are purely random events, then an accident is a premature truncation in the intended lifespan of a rental vehicle and serves as a *de facto* randomized experiment that can be used to infer the effect of adopting a shorter replacement threshold on profitability. The average internal rate of return on the 123 vehicles that were scrapped due to accidents in the data set was 36%, which is statistically significantly below the 47% average internal rate of return for all cars in the data set. However there are reasons to be suspicious of accidents as an instrumental variable: if there are "lemons" problems (unobserved problems with certain

cars) that make certain vehicles both less profitable and less safe to drive, then the occurrence of an accident could be an indicator of a lemon, and not necessarily evidence that replacing cars earlier reduces profits. Personally, we find this argument a bit far-fetched, but the other major limitation of the instrumental variable approach is something we find more compelling: the accident “experiments” only tell us the effects on profitability of replacing cars too early, *but it does not tell us how profits would change if the company kept its cars longer than it currently does.*

Thus, there are only two other remaining possibilities for how we might go about testing the hypothesis that this company is a profit maximizer. One is to undertake one or more *controlled experiments*, that is, to pick one or more car types at one or more of the company’s locations, and randomly assign some cars to the *treatment group*, where the “treatment” would correspond to keeping cars longer than the company currently does, and cars in the *control group* would continue to be subject to the company’s existing replacement policy. By following the cars in the treatment and control group for a sufficient length of time (i.e. from their initial purchase until they are sold), we can compare their profits/returns. If the cars in the treatment group have higher average profits or returns, this would constitute evidence against the hypothesis that the company’s existing operating policy is optimal (i.e. profit maximizing).

The drawback of controlled experiments is that they are costly and time-consuming. Further there are many possible “treatments” that one could imagine testing: the treatments could involve various combinations of keeping cars longer for various durations in terms of age or odometer relative to the *status quo* as well as various choices about how to discount rental rates for older vehicles. We need sufficiently many vehicles in the treatment and control group to make statistically significant inferences, so the number of possible experiments that the firm could undertake at any point in time is strictly limited. For these reasons, it appears that an experimental approach to testing whether the company is profit-maximizing is not very promising.

The only remaining approach (at least of which we are aware) is to construct an *econometric model* of the firm’s operations. This model can be *simulated* to generate predicted outcomes both under the *status quo* and under a variety of alternative hypothetical replacement and operating strategies. The key advantage of the modeling/simulation approach is that simulations are very cheap, and a large number of alternative scenarios and operating strategies can be evaluated extremely rapidly. The key limitation to this approach is that if the econometric model does not provide a good approximation to the actual operations of this company, its predictions of the effects on the firm’s profits from implementing various hypothetical alternative operating strategies will not be trustworthy.

We will adopt the modeling/simulation approach in this paper. In the next section we present the econometric model of the company’s operations, and in the section after that we simulate the model and show that it provides a good approximation to the actual outcomes for this company under its *status quo* operating strategy. Thus, we argue that the modeling/simulation approach is trustworthy, although we still recommend that the predictions of the model be *validated* by conducting a controlled experiment to evaluate whether the optimal replacement strategy implied by this model really does lead to the significant increase in profits that the model predicts.

3 An Econometric Model

In order to get more insights into the behavior of the rental company and to evaluate the profitability of its vehicle replacement decisions, this section describes an econometric model of the company’s vehicle rental operations. We introduce a *semi-Markov model* in which cars that the company owns can be in one of four possible states at any given time: 1) in a long term rental contract (i.e. a “long term rental spell”), 2) in a short term rental contract (i.e. a “short term rental spell”), 3) in the lot waiting to be rented when the previous rental state was a long term rental spell, and 4) in the lot waiting to be rented when previous rental state was a short term rental spell. We refer to the latter two states, 3 and 4, as *lot spells*. We differentiate between these states since it turns out empirically that the duration distribution of a car in a lot spell is quite different depending upon whether it had previously been in a long or short term rental contract.

A *semi-Markov process* is a stochastic process that can be in one of a finite number of possible states at any given time, but where the duration distributions in each of these states (also called the holding time distributions) can be arbitrary distributions. A *Markov process* is a special case of a semi-Markov process where the duration distributions in each state are restricted to be exponential (or geometric, in the case of discrete time models). In this study we use a discrete time model, with the relevant time unit being one day. Let r_t denote the rental state of a given car on day t . From the discussion above, r_t can assume one of the four possible values $\{1, 2, 3, 4\}$.⁴

In addition to the rental state, other relevant state variables for modeling the decisions of the rental company are the vehicle’s *odometer value*, which we denote by o_t , and the *duration in the current rental*

⁴Actually we could distinguish a fifth possible state, $r_t = 5$, denoting a brand new car that is in its first lot spell. Empirically we have found that the duration distributions for the initial lot spell can be well approximated as a mixture of the duration distributions for lot spells $r_t = 3$ and $r_t = 4$, and so to reduce the size of the state space, we use only four possible values for r_t and probabilistically assign new cars to lot states 3 and 4 in such a way that the initial duration distribution closely matches the distribution of initial lot spells that we observe in the data.

state, which we denote by d_t . Thus, we seek to model the joint stochastic process $\{r_t, o_t, d_t\}$. There is another potential state variable of interest, the vehicle's *age* which we denote by a_t . If we let $t = 0$ denote the date at which a car was bought, and if $a_0 = 0$ (when a car is acquired, it is a brand new car), then we have $a_t = t$, i.e. the age of the car in days is the same as the time index t .

In the empirical analysis below, it turns out that a vehicle's age t is strongly correlated with its odometer value o_t . Because of this "collinearity problem" it is difficult to identify the independent effects of these two variables on decisions to sell a car, or on maintenance costs, state transition probabilities, durations in states, and even on the resale price of used vehicles. Since there are numerical and computational advantages to minimizing the number of different variables we include in the dynamic programming model, we have opted to exclude vehicle age from the list of variables that we use to predict the company's selling decision, vehicle resale prices, transitions and durations in spells, and so forth. However, the model does in fact keep track of vehicle age, and as we demonstrate in section 4, simulations of the model accurately predict the mean age at which the company sells vehicles, even when we restrict the model of the company's sales policy to depend only on the vehicle's odometer o_t and not its age a_t .

Using the three key variables $\{r_t, o_t, d_t\}$ we will also be able to simulate realizations of *rental revenues* and also *maintenance costs* using the data that the company provided us. With this information, we can construct a complete econometric model of the company's rental operations, and conduct stochastic simulations of the model to see how accurately it can represent the company's actual operations. The econometric model requires us to specify and estimate the following objects: 1) a model of the resale price the company receives if it were to sell one of its cars, 2) a model of the random durations of a car in each of the rental and lot states, 3) a model of a car's transition to the next rental spell at the end of the current rental or lot spell, 4) a model of the utilization (kilometers driven) on a particular car during a long or short term rental contract, 5) a model of rental revenues received and maintenance costs incurred by the company over the life of the car, and 6) a model of the company's *selling decision*.

We will now discuss each of these components in turn, describing the econometric model and our empirical findings. The first model is a regression equation to predict the resale price of a rental car when it is sold. We have data on both the new price $\bar{P}(\tau)$ as well as the realized sales price $P_t(o_t, \tau)$ of each car, where τ denotes a particular make and model of vehicle, which we will also call a *car type*. The econometric analysis will focus on the three car types discussed in section 2, i.e. a particular make/model of compact, luxury, and RV. We wish to emphasize that in order to maintain confidentiality of the data, we are not able to disclose the specific brand and model of these three car types, and instead use the rather

vague car type designations to refer to them. Thus, *whenever we refer to one of these car types, such as “compact”, we are not referring to the class of all compact cars owned by this company, but instead to a specific brand and model of compact.*

For each of the three car types τ , we estimated a simple linear regression model with the logarithm of the depreciation rate, $\bar{P}_t(o_t, \tau)/\bar{P}(\tau)$, as the dependent variable

$$\log\left(\frac{P_t(o_t, \tau)}{\bar{P}(\tau)}\right) = \alpha_1(\tau) + \alpha_2(\tau)o_t + \varepsilon_t. \quad (3)$$

The results from this model can be interpreted as a regression with type-specific “depreciation coefficients” $(\alpha_1(\tau), \alpha_2(\tau))$ where α_2 measures the effect of the car’s odometer value on its resale price. We also estimated regressions where we included the vehicle age and other variables, such as the number of accidents and the total accident repair cost as predictors of the resale price of a car. These results are presented in Table 1 below.

Variable	Compact All Locations		Luxury All Locations		RV All Locations	
	Estimate	(t-stat)	Estimate	(t-stat)	Estimate	(t-stat)
Constant	-0.4789	(-7.61)	-0.6201	(-20.85)	-0.8521	(-4.04)
Age (days)	-0.0001	(-2.53)	-0.0004	(-5.67)	-0.0004	(-2.17)
Odometer (000 km)	-0.0007	(-2.11)	-0.0011	(-2.10)	0.0016	(1.91)
Number of Accidents	-0.0112	(-1.05)	0.0006	(0.04)	0.0371	(1.00)
Accident Repair Costs	$-0.8.88e^{-6}$	(-1.04)	$-4.672e^{-6}$	(-0.57)	$-1.654e^{-6}$	(-0.56)
Internal Rate of Return	0.1629	(12.21)	0.067	(0.99)	0.394	(4.43)
Maintenance Cost per Day	0.0092	(0.64)	-0.0039	(-0.31)	-0.0053	(-0.33)
N, R^2	288	0.389	91	0.420	41	0.481

Table 1: Regression Results for Dependent Variable $\log(P_t(o_t, \tau)/\bar{P}(\tau))$

The constant term in the regressions is a measure of how much depreciation a vehicle experiences the “minute it goes off of the new car lot.” We see that this predicted “instantaneous depreciation” is huge for all three vehicle types, but is significantly lower for the compact, which retains 62% of its original value ($= \exp(-.48)$) the instant it is driven off the lot, compared to only 52% for the luxury vehicle and 43% for the RV. Figure 2 provides a scatter plot of the resale prices for the compacts, graphed against their odometer value at time of sale (the results for the luxury and RV are similar but not shown due to space constraints). The rapid early depreciation in car prices is evident in these graphs. While a number of cars are sold quite “early” after their initial purchase (measured either in terms of their age or odometer value), we do not have any observations of sale prices the company might have received if it were to have sold vehicles in

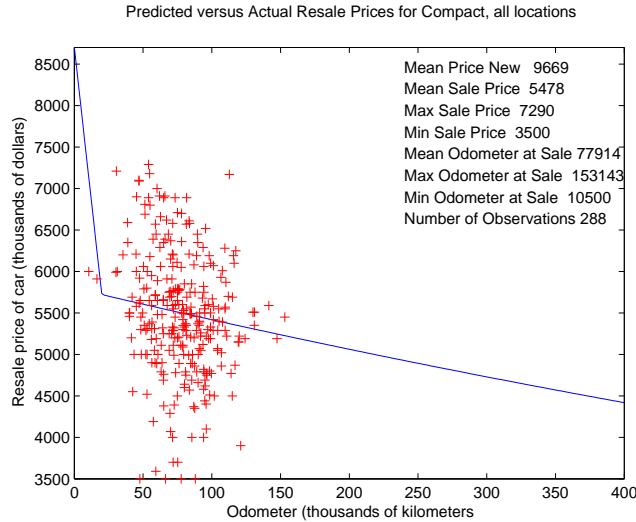


Figure 2 Predicted versus Actual Resale Prices: Compact – All Locations

only a matter of a few weeks or months after the initial purchase. For the purposes of our modeling, we did not feel we could trust the regression extrapolations for used vehicle prices for age or odometer values very close to zero. Therefore, we made a simple linear extrapolation to predict the prices that very new used cars would sell for (i.e. used cars with fewer than 10,000 miles). Figure 2 show that even for relatively low odometer values the regression accurately predicts the mean resale price. It is immaterial whether we use a straightline extrapolation or assume that the resale function has a discontinuity at zero, since the company almost never sells cars that have fewer than several thousand kilometers on their odometers.

Next we consider the econometric estimation vehicle usage during rental spells. The company does not have mileage charges, and places no constraints on its customers’ choice of how far to drive during their rental contracts. The intensity of utilization by rental customers is obviously an important consideration because it determines how quickly a car will “age” in terms of its odometer value, and the odometer is in turn a key predictor of the resale value of the car. However the difficulty, noted above, is that the firm frequently does not accurately record the in and out odometer values for its vehicles, making it hard for us to determine how far a car was driven on particular rental spells. To get around this problem and make inferences about the conditional probability distribution of the number of kilometers driven of a rental contract of type $r \in 1, 2$ and duration d , we need to impose some additional assumptions.

Let $F(o'|o, d, r)$ denote the conditional distribution of the (frequently unobserved) odometer value on a rental car that has returned from a rental contract of type r , lasting d days, when the out odometer value was o (i.e. the car had an odometer reading of o at the start of the rental spell). Thus, $\nabla o = o' - o$ is

the number of kilometers driven by the customer during the rental spell. We assume that the number of kilometers traveled each day by a rental customer are *IID* (Independent and Identically Distributed) draws from an exponential distribution with parameter λ_r . Conditional on spell length d , it follows that $F(o'|d, r)$ is a gamma distribution, since a sum of *IID* exponential random variables has a gamma distribution.⁵ The probability density function corresponding to F is given by

$$f(o'|o, d, r) = \begin{cases} \frac{[o'-o]^{d-1} \exp\{-(o'-o)/\lambda_r\}}{[\lambda_r]^d \Gamma(d, \lambda_r)} & \text{if } o' - o \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where $\Gamma(d)$ is the *Gamma function*. Thus, $o' - o$ is the actual number of kilometers traveled by the customer during the rental contract. We have $E\{o' - o|o, d, r\} = d\lambda_r$, so we can interpret λ_r as the mean number of kilometers traveled per day in a rental contract of type r . For notational consistency, we set $\lambda_r = 0$ if $r > 2$, i.e. cars do not travel any kilometers when they are on the lot waiting to be rented.

In order to estimate kilometers traveled per day under short term and long term contracts, it would be natural to look to the rental contract data directly and take the average kilometers traveled per day for short term and long term contracts separately. Since the in and out odometer values for rental contracts do not appear to be accurately recorded in the company's data, we cannot use this approach. If we had such data, it might even be possible for us to estimate the conditional distributions $F(o'|o, d, r)$ non-parametrically or semi-parametrically, and thus, not have to rely on the parametric assumption that kilometers traveled per each day of a rental contract are *IID* exponential variables. In the absence of contract-specific mileage data, we can still estimate the two λ_r parameters necessary for us to determine the distributions $F(o'|o, d, r)$ using the accurate records we do have on the odometer value of each vehicle at time of sale.

Suppose that at time of sale, a rental car had been rented for N^s days under short term rental contracts and N^l days under long term rental contracts. Then the odometer value on the car is given by

$$\tilde{o} = \sum_{i=1}^{N^l} \nabla o_i^l + \sum_{i=1}^{N^s} \nabla o_i^s \quad (5)$$

where ∇o_i^l and ∇o_i^s are the realized number of kilometers traveled under long and short term contracts, respectively. Under our assumptions that kilometers traveled per day are exponential random variables with parameters λ_1 (for long term contracts) and λ_2 (for short term contracts), we have

$$E\{\tilde{o}|N^l, N^s\} = \lambda_1 N^l + \lambda_2 N^s. \quad (6)$$

⁵Actually, the distribution is part of a special subclass of the Gamma family known in renewal theory as the *Erlang distribution* since the parameter α of the Gamma distribution is an integer $\alpha = d$.

Since we do observe the odometer at sale \tilde{o} and number of days a vehicle is rented, N^l and N^s , this implies that we can estimate λ_1 and λ_2 as coefficients on a simple linear regression

$$o_i = \lambda_1 N_i^l + \lambda_2 N_i^s + \varepsilon_i \tag{7}$$

where o_i is the odometer at time of sale on the i^{th} rental car sold by the company, and N_i^s and N_i^l are the number of days the i^{th} car had been in short and long term rentals over its service life.

We use the regression estimates and the information on (N_i^l, N_i^s) for each car, now indexed by the day in its service life, t , to compute a predicted value for the car’s odometer, $\hat{o}_t = \hat{\lambda}_1 N_t^l + \hat{\lambda}_2 N_t^s$, at day t in the car’s life. The high R^2 values for the odometer regressions in equation (7) give us confidence that the predicted odometer values are reasonably accurate.⁶ With \hat{o}_t , the next step is to analyze the determinants of the company’s decision to sell its cars.

Variable	Compact	Luxury	RV
	All Locations	All Locations	All Locations
λ_1	78.7	86.6	95.4
λ_2	157.1	140.8	167.7

Table 2: OLS Estimates of λ_1 and λ_2 in the regression equation (7)

We estimated a reduce-form binary logit model to capture the company’s *status quo* replacement policy. We do not report the detailed results here: it suffices to note that the only significant predictors are the age and odometer of the vehicle. Figure 3 summarizes the firm’s replacement policy for compact cars (the replacement policy for luxury and compact is very similar). It shows the cumulative distribution for replacements as a function of the odometer value (left hand panel), and the cumulative distribution of the vehicle age (right hand panel). The right hand column confirms the company’s claim that the target replacement age for its vehicles is 3 years. The mean age of the three types of cars at replacement is fairly close to this three year target: 2.8, 2.9 and 2.7 years for the compact, luxury and RV, respectively. However the left hand panel shows that in terms of odometer values at replacement, there is greater variability. The mean odometer value at replacement for the three vehicle types is 78, 75 and 89 thousand kilometers, respectively. The fact that mean replacement ages vary much less across the three car types than the mean odometer values at replacement may be taken as evidence that the company bases its replacement decision

⁶We also tested the accuracy of the regression predictions by comparing the actual odometer reading that was recorded at dates where maintenance was performed with the predicted values \hat{o} on these same dates. The distribution of prediction errors is centered at zero with small variance, indicating our regression model is a very good predictor of actual vehicle usage.

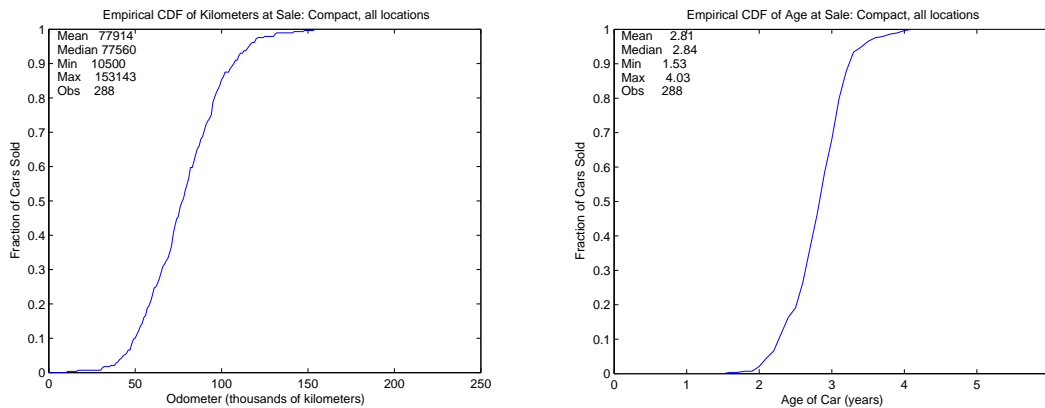


Figure 3 Cumulative Fractions of Cars Replaced: Compact – All Locations

more on the “three year rule” than on a rule based on number of kilometers driven. As we will see in the next section, due to the high degree of collinearity between age and odometer values, a replacement rule based on odometer value can provide a good approximation to an age-based replacement rule and vice versa.

The remaining objects to be estimated are the *spell durations* and the *spell transition probabilities*. As is well known, there is a duality between duration distributions and the corresponding *hazard functions*. We chose to work with hazard functions and let $h(d, r)$ denote the hazard rate for the rental state r , i.e. it is the conditional probability that the car that has been in rental state r for d days will exit the state d on the next day, $d + 1$, and with the complementary probability $1 - h(d, r)$ it will continue to remain in state r . Since we have sufficiently many observations of rental spells, we were able to estimate the hazard functions for these spells non-parametrically. The longest duration for any rental spell is 31 days, i.e. the maximum duration of a monthly rental. As one might expect, the duration distributions of long and short term rental spells are very different: most short term rentals last only a few days, whereas most long term rentals last for an entire month. There is only minor variation in the durations of long term rentals, i.e. some rentals are for 29 days, 30, or 31 days. Most long term rentals last more than 15 days, perhaps in part due to the 20% penalty the company imposes on early return of vehicles in a long term contract.

We have far fewer observations on lot spell durations, especially for type 3 lot spells, i.e. where the previous rental spell was a long term contract. This is due to the high probability of roll overs in longer term contracts, leading to relatively few observations on intervening lot spells with positive durations. Due to the relatively small number of observations, the nonparametrically estimated hazard functions are quite jagged. Also, unlike rental contracts, there is no *a priori* upper bound on the duration of a lot spell. As

a result we needed some method of extrapolation to predict durations given that we have only a small number of cases with extremely long lot durations. Our solution to this problem was to assume that the hazard function is constant after $d = 31$ days, which implies a geometric upper tail for the distribution of lot spells. We estimated this constant upper tails of the hazard function by imposing the constraint that the implied duration distribution (with a smoothed, non-parametrically estimated lower tail and the geometric upper tail) has a mean duration that equals the actual mean duration for type 3 and 4 lot spells, respectively.

When a spell in a given rental state ends, there is a transition to a new rental state. Let $\pi(r'|r, d, o)$ denote the probability that the new rental state for a car will be r' given that the current rental state is r , the odometer value is o , and the duration in state r is d . We call π the *rental state transition probability*. Recall that if $r > 2$, the car is in a lot spell. We rule out the possibility of “self transitions” to the lot (that is, we assume that $\pi(r|r, d, 0) = 0$ for $r > 2$) because the hazard function $h(d, r)$ already provides the probability that the lot spell has ended, and there is no conceptual difference between a lot spell continuing for one more day, versus the case where a lot spell terminates and immediately re-enters the lot via a self-transition $r' = r$. Thus the restriction $\pi(r|r, d, o) = 0$ for $r > 2$ can be viewed as an econometric identification normalization.

However, for rental spells, there is a conceptual distinction between a rental spell that terminates with an immediate transition to a new rental spell versus the case where an existing rental contract continues for one more day. The former case can be viewed as an immediate roll over of one rental contract to another one, such as when a previous customer renews or extends their previous rental contract by one more month (in the case of a long term contract), or by another day (in the case of a short term contract). Thus, we allow $\pi(r|r, d, o) > 0$ for $r \in \{1, 2\}$, and interpret this probability as a probability of a contract extension or roll over.

The rental spell transition probability can also accommodate transitions from a rental spell to a lot spell, except that by our definition of the two types of lot spells, it must be the case that $\pi(4|1, o, d) = 0$ and $\pi(3|2, o, d) = 0$, i.e. if a car is leaving a long term rental spell, it can only transition into a lot spell of type 3 which is defined as a lot spell where the previous rental spell was a long term contract. Similarly, a car leaving a short term rental spell can only transition into a lot spell of type 4. The reason why we distinguish the two types of lot spells is that the hazard functions and mean durations for type 3 lot spells are different than for type 4 lot spells. In particular, for all three types of cars, hazard rates for type 3 lot spells are lower and mean durations are higher. In plain language, if a car had previously been in a long term rental that did not immediately roll over, one would expect the car to be on the lot for a longer period

of time compared to the case where the car has returned to the lot from a previous short term rental.

Since there are three possible destination states for transitions out of a rental spell i.e. immediate return to a long term or short term rental, or back to the lot, we used a trinomial logit model to estimate these probabilities. This probability is given by

$$\pi(r'|r, d, o) = \frac{\exp\{v(r, d, o)\theta_{r'}\}}{\sum_{\rho \in \{1, 2, l(r)\}} \exp\{v(r, d, o)\theta_{\rho}\}}, \quad (8)$$

where $v(r, d, o)$ is a vector-valued function of the variables (r, d, o) and θ_{ρ} is an alternative-specific vector of parameters, for $\rho = \{1, 2, l(r)\}$ (where $l(r)$ denotes a lot spell, either of type 3 if $r = 1$ or type 4 if $r = 2$) with the same dimension as v . As is well known, it is not possible to identify all three of the θ_{ρ} vectors. Therefore we make an identifying normalization that $\theta_1 = 0$, i.e. we normalized the parameters for transition to long term contract to zero.

Since our identification normalization rules out “self-transitions,” there are only two possible destinations from a lot spell: i.e. to either a long term or short term rental spell. We used a binary logit model to estimate transition probabilities out of lot spells. Similar to the trinomial logit specification, we used transition probabilities specified as

$$\pi(r' = 1|r, d, o) = \frac{\exp\{v(o, d)\theta_r\}}{1 + \exp\{v(o, d)\theta_r\}}, \quad r \in \{3, 4\}. \quad (9)$$

We do not report the actual estimation results here, but there are two main findings we think are noteworthy:

- 1) for all car types, there is a very high probability that cars will be initially rented in long term contracts,
- 2) the results provide clear evidence of “contract age effects”. That is, as the odometer value increases (i.e. the age of the car increases) the probability of transitions into long term rental contracts decreases, the probability of transitions into short term rental contracts increases, and the probability of remaining on the lot unrented increases.

The remaining objects in the econometric model of this company are the daily rental rates and maintenance costs. Daily rental rates do not need to be estimated since we can simply access the company’s published tariff rate. However there is a small amount of variability in daily rental rates due to variations in optional equipment and features on cars (e.g. some cars have larger engines than the standard size, etc.). To account for this variation, we simply computed the mean daily rental rate by dividing the total rental revenues earned in long and short term rental spells by the number of days in these spells.

Maintenance costs are incurred on an episodic basis. The company appears to adhere to a fairly regular periodic maintenance schedule with operations such as oil changes, brake pad replacement and so forth

occurring at regular intervals (either intervals of time such as every 3 months, or in terms of odometer values, such as every 50,000 kilometers, etc). However, there is also evidence of a high frequency of “unexpected maintenance” resulting from random malfunctions or problems in particular cars. While we could have tried to model the durations between successive maintenance visits, with a conditional distribution of maintenance costs that are incurred at each maintenance event, we opted for a simpler approach that appears to work just as well. Our approach is simply to charge a daily equivalent maintenance rate, where we estimate the daily maintenance charge by taking the mean of the ratios of total maintenance costs over the service life of the vehicle divided by the age of the vehicle at the time it was sold. In the next section we will now show, via the stochastic simulations, that this simplified treatment of maintenance costs does not compromise our ability to provide a good overall model of the firm’s operations.

The final issue we address is whether there is evidence of aging effects in vehicle maintenance costs. Scatter plots of the average daily maintenance costs incurred by the company on its vehicles as a function of the predicted odometer value show no evidence that daily maintenance costs increase with the odometer value of the vehicle, at least over the range of odometer values observed in this company’s fleet. Thus, we conclude that the only aging effects that we can detect in the econometric analysis are: 1) the rapid decline in resale values of vehicles as a function of their age and odometer value, and 2) the “rental composition aging effect”, i.e. the tendency for cars to be initially rented on long term contracts, but to gradually transition to an increasing share of short term rental contracts and to spend more time on the lot as the vehicle ages.

4 Evaluating the Econometric Model: Simulated versus Actual Outcomes

In the previous section we described and estimated an econometric model of the rental company’s operations. In order to determine if this is a good model that accurately captures the key features of the behavior of this company, this section presents comparisons of simulated outcomes from the econometric model to the actual outcomes for each of the three vehicle types analyzed in section 3.

Our approach to simulating the econometric model is conceptually straightforward. Starting with a new car on the lot with an odometer value of zero, we simulate arrival of customers using the estimated transition probabilities discussed in section 3. Once a car is rented (either in a long or short term rental spell), we use the estimated hazard functions to determine the duration of the rental spell, and the estimated regression coefficients (from the regression equation (7) that determines average daily kilometers traveled

in short and long term rental contracts, respectively) to predict the number of kilometers driven during the rental spell. More precisely we use the fact that under the assumption that daily kilometers traveled are *IID* draws from an exponential distribution with parameter λ_1 (for short term rentals) or λ_2 (for long term rentals), and the total duration of the rental spell to generate a random draw of the total number of kilometers driven during that spell. We continue this process until the first time a simulated car re-enters the lot with an odometer reading o_t that exceeds its *replacement threshold* \bar{o} . If $o_t > \bar{o}$, the car is sold (where we use a draw from a lognormal distribution determined by the used car sale price regression to predict its sales proceeds), a new car is purchased for a price $\bar{P}(\tau)$, and the life of the new simulated car begins. Finally, we draw the replacement thresholds \bar{o} at the start of the life of each newly purchased car, using the empirical distribution of odometer values at which the company actually replaces its cars in the left hand panel of Figure 3.

Figure 4 presents comparisons between simulated and actual distributions of the odometer and vehicle age (in days) at which vehicles are replaced. We present the results for the luxury car type only, although the results for the compact and RV are similar. The left hand panel of Figure 4 compares the actual distribution of odometer values (solid blue line) with the simulated distribution (dashed red line). We see that the two distributions are close to each other, which is a result we would expect given that we have drawn the odometer thresholds that determine when vehicles are to be replaced in the simulations from the actual (empirical) distribution. Thus, the differences in the two distributions in the left hand panel of Figure 4 are entirely due to *sampling error* in our random sample of 100 simulated cars.

The right hand panel of Figure 4 compares the actual distribution of replacement ages to the one implied by the econometric model. In this case we do not directly draw the age at which a car is replaced from the empirical distribution of replacement ages, so there is no guarantee that the simulated distribution of ages at replacement is close to the actual distribution. Indeed, the simulated age at replacement is a result of a more complicated set of interactions that depend on other estimated objects in the econometric model that determine the number of times a vehicle was rented, the durations of these rental spells, and the numbers of kilometers driven per rental spell. This implies a particular co-evolution of vehicle age and odometer values, so that when the simulated odometer value exceeds the replacement threshold \bar{o} that we randomly drew from the empirical distribution in the left panel of Figure 4 at the start of each car's simulated history, the random time at which the car's simulated odometer exceeds \bar{o} determines the simulated lifetime (age) of the vehicle in question.

The simulated and actual distributions of ages at replacement are further apart than the distributions of

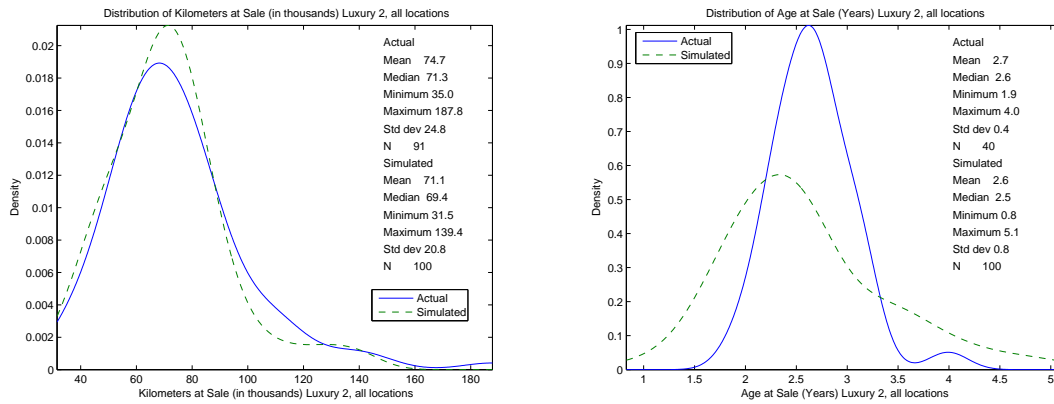


Figure 4 Simulated versus Actual Odometer and Age at Sale: Luxury – All Locations

odometer values at replacement, although we note that the mean simulated age at replacement, 2.6 years, is very close to the actual value, 2.7 years. In particular, the simulated distribution of replacement ages has a larger variance, with more replacements at younger ages and also at older ages compared to the actual distribution. This discrepancy probably reflects the fact that the company’s replacement decisions are based more on age than odometer value, and thus actual replacements are more tightly concentrated around the three year replacement target that we discussed in section 3. Even though age and odometer values are highly correlated with each other, a purely odometer-based approximation to the company’s replacement rule can be expected to result in a larger variation in replacement ages. Thus, cars that have high simulated capacity utilization and large numbers of simulated rentals will be younger than average at time of replacement, whereas those with low simulated utilization rates will be older than average at time of replacement.

We did not adopt a more complex replacement rule based both on age and odometer and other variables such as duration in the lot spell before the vehicle was replaced because we feel that the simpler odometer-based replacement rule provides a sufficiently good approximation to the company’s behavior and outcomes — as we verified by comparing simulated versus actual distributions of outcomes for twelve different outcome variables of interest. Another reason motivating our use of an odometer-based approximation to the company’s *status quo* replacement policy is that this enables us to adopt a stationary Markovian decision process formulation to estimate the expected present discounted value of the company’s profits over an infinite horizon. We argue that the infinite horizon benchmark (which values the discounted profits from an infinite sequence of rental cars, not just the currently operating rental vehicle) provides a more reasonable basis for comparing the profitability of alternative operating strategies than a

finite horizon benchmark, which calculates the discounted profits earned only by the current generation of autos over their lifetimes.

Table 3 compares actual versus simulated values of a number of different moments of the various quantities of interest, including the number of long term and short term rental spells, the number of lot spells, and the total number of days spent in each of these spells. The model not only does a good job of matching the mean number of spells and durations of each spell type, but it also does a good job of capturing the overall distribution as well. Matching both the number and duration of the various types of spells turns out to be the key to accurate predictions of revenues, profits, and returns.

Table 3 also compares simulated versus actual moments of the relevant financial variables, including the mean proceeds from sales of cars, mean maintenance costs, and mean short and long term revenues. The fact that simulations of the model match actual mean proceeds from sales of vehicles suggests that the simple lognormal regression model of vehicle price depreciation is a good one. Similarly, we find that the model accurately predicts mean maintenance costs, although the simulated distribution of maintenance costs has lower variance than the actual distribution. This is to be expected given our shortcut procedure for simulating maintenance costs described in the previous section. That is, instead of trying to estimate the durations between successive maintenance episodes, we predicted cumulative maintenance costs using the age of the car multiplied by the estimated equivalent daily maintenance cost. Thus, total maintenance costs simply equal the service life of the car (in days) times the average maintenance cost expenditures per day. While it is certainly possible to improve on the way we model maintenance costs, maintenance costs are a distinctly second order aspect of the rental car business, in the sense that these costs are dwarfed by rental revenues and the purchase and resale price of the vehicle. Indeed, we see that the total maintenance cost is only about 1/12th of the average resale value for the luxury car type, and about 1/23rd of the cost of a new car.

Table 3 compares simulated and actual mean values of a number of quantities of interest for the luxury vehicle (results of a comparison of the simulated and actual values for the compact and RV lead to similar qualitative conclusions). If we plotted the entire distributions of outcomes for each of these quantities, the reader would see that the model does a good job of capturing the overall distribution of outcomes and not just the mean values. For example, the model accurately predicts that long term contracts account for over 80% of the rental revenues earned. This is not surprising given that of the approximately 800 days these cars were rented on average over their 985 day service life, nearly 90% of the rental days were in long term contracts. However, it may seem surprising in view of our finding that there is a negative relationship

Variable	Actual Mean (standard error)	Simulated Mean (standard error)
Number of long term rental spells	24.7 (9.3)	22.8 (7.0)
Days in long term rental spells	697 (271)	643 (213)
Number of short term rental spells	15.2 (18.7)	18.4 (18.3)
Days in short term rental spells	100 (124)	101 (104)
Number of lot spells	13.9 (14.2)	14.2 (12.8)
Total days on the lot	193 (141)	129 (110)
Car sales proceeds	12,283 (1,667)	12,109 (1,710)
Total maintenance costs	956 (932)	1,086 (360)
Revenue from long term rentals	28,207 (10,632)	28,622 (9,606)
Revenue from short term rentals	6,106 (7,337)	7,446 (7,588)
Total (undiscounted) profits	22,244 (6,224)	22,212.7 (9,608)
Internal rate of return	49.2% (10.4)	47.1% (6.9)
Number of observations	40	100

Table 3: Comparison of Actual versus Simulated Outcomes — Luxury

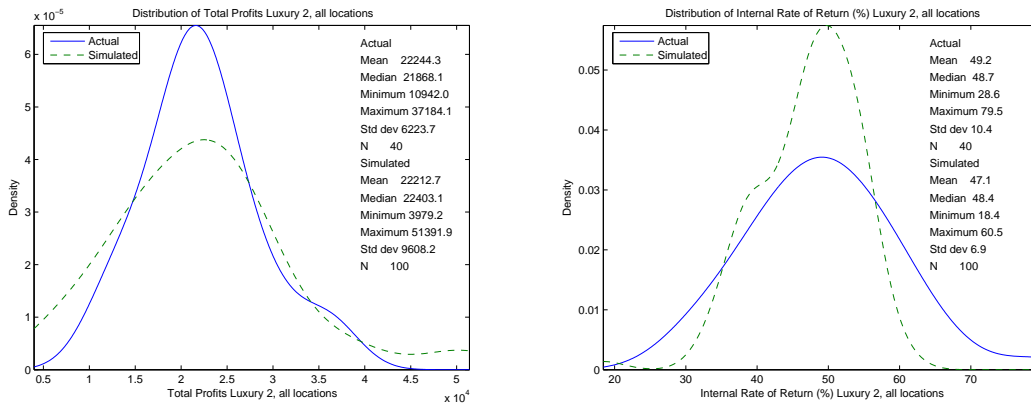


Figure 5 Simulated versus Actual Internal Rates of Return: Luxury – All Locations

between the fraction of time spent in long term contracts and the IRR on the vehicle. This regression estimate suggests that long term contracts are less profitable than short term contracts. However, as we noted above, there are reasons to distrust any conclusions based on the simple IRR regression in Table 1. See Cho and Rust (2007) for further discussion of the relative profitability of long versus short term rental contracts.

Figure 5 compares the actual and simulated distributions of total profits and internal rates of return, IRRs, for the luxury vehicle type (results for the compact and RV car types are similar). Once again the econometric model provides a good prediction of mean total profits and mean IRR, although the simulation results show that the distribution of simulated total profits has a larger variance than the actual distribution,

whereas distribution of simulated IRRs has a lower variance than the the actual distribution. We are not quite sure why the econometric model should overpredict the variance of total profits and underpredict the variance of IRRs. However, the discrepancies might have significant implications for the predictions of the effect of different strategies on expected discounted profits if the company is an expected profit maximizer (i.e. the company is not risk averse). In that case only the mean values of profits matters and an expected profit maximizing firm would be indifferent between two different operating strategies that result in the same mean profits, even though one of the strategies results in a larger variance of profits.

While we certainly do not claim that this is a perfect model of all aspects of the company's operations (we have not presented Chi-square or Kolmogorov-Smirnov goodness of fit statistics since we expect that the model would be formally rejected by such specification test statistics), it does seem to provide a reasonable first approximation to the company's operations. Note that we did not estimate the econometric model with the goal of trying to minimize the distance between simulated and actual outcomes. Instead, we estimated each component of the econometric model separately, and evaluated the implications of the overall model via simulations. If we had estimated all of the parameters of the model jointly in order to maximize a likelihood function or a simulated minimum distance (or method of moments) criterion, we could no doubt produce simulations that result in an even closer fit to the data. However, even in this case, we believe it is likely that formal goodness of fit or specification test statistics would reject the model, as typically happens for most parametric econometric models when there are a sufficiently large number of observations. We believe that econometric models can still have credibility even if the models are rejected by formal specification tests. The relevant question is whether there are alternative models that fit significantly better. We are not aware of alternative econometric models that could do a better job of fitting the data than the model we estimated in the previous section while still retaining the tractability and flexibility to make the discounted profit calculations we present in the next section. In the remainder of this paper, we will focus on the question of trying to determine what new insights can be derived in the typical case where we make a serious effort to find an econometric model that provides a reasonable, but by no means perfect, approximation to the actual data generating mechanism.

In any case, from our perspective, the econometric model we have formulated provides a sufficiently good approximation to the company's actual operations that we think it should be a credible model to use to evaluate the consequences of *certain* modifications in the company's operating strategy. That is, we can simulate the econometric model under a range of alternative hypothetical scenarios, and use it to predict profits and rates of return and see how these compare to the company's *status quo* operating

policy. However for reasons we will elaborate on shortly, there are certain modifications to the company’s operating strategy for which we have little data available to base a prediction. An example would be the predicted effect of a large increase in rental rates. Of course, we would expect a large rise in rental rates would lead to fewer rentals, and this would change the stochastic structure of durations and transitions between rental states. Since we do not have any observations on large variations in rental rates in the past, we have no basis for estimating or extrapolating how the stochastic structure of the econometric model, and thus the implied distribution of profits would change as a result of significant increases or decreases in rental rates. Thus, we need to exercise caution and clearly demarcate hypothetical simulations for which we lack adequate data to make a reliable prediction about how certain changes in the company’s operating strategy would affect its expected revenues and profits.

5 Simulating Profitability of Counterfactual Replacement Policies

While it is possible to evaluate *specific hypothetical alternatives* to the company’s *status quo* operating policy using simulation methods similar to the previous section, there are more efficient methods available for characterizing the *optimal replacement policy* that involve searching over what is effectively an infinite dimensional space of *all possible replacement policies*. Mathematically, the optimal replacement problem is equivalent to a specific type of *optimal stopping problem* known as a *regenerative optimal stopping problem* (see Rust, 1987). The term “regenerative” is used, since the decision to replace a vehicle does not stop or end the decision process, but rather results in a “regeneration” or “rebirth”, i.e. a replacement of an old vehicle by a brand new one.

Cho and Rust (2007) use numerical dynamic programming to formulate and solve the optimal stopping problem. The optimal replacement policy takes the form of a *threshold rule*, i.e. the optimal time to replace a car occurs when its odometer value o exceeds a threshold value $\bar{o}(d, r, \tau)$ that depends on the current rental state r , the duration in that state d , and the car type τ . Using numerical methods, they solved the dynamic programming problem and calculate the optimal stopping thresholds $\bar{o}(d, r, \tau)$ for the compact, luxury and RV car types and the associated optimal *value functions* $V(r, d, o, \tau)$. The value function provides the expected discounted profits (over an infinite horizon) under the optimal replacement policy for a vehicle of type τ that is in state (r, d, o) .

It is also possible to compute the value of any alternative operating strategy μ , which can include *mixed* or probabilistic operating strategies where the decision to replace a car is given by a conditional

probability distribution $\mu(r, d, o, \tau)$. Let $V_\mu(r, d, o, \tau)$ denote the expected discounted profits (again over an infinite horizon) under the alternative replacement policy μ . We will calculate both V and V_μ where μ is an approximation to the company's *status quo* operating policy. Thus, the difference $V(r, d, o, \tau) - V_\mu(r, d, o, \tau)$ will represent our estimate of the gain in profits from adopting an optimal replacement policy. As we noted in the introduction, the optimal policy entails keeping cars significantly longer than the company currently keeps them, but by doing this we show that the company can significantly increase its expected discounted profits.

If we were to solve the regenerative optimal stopping problem under the assumption that the only aging effects are: 1) the depreciation in vehicle resale values, and 2) the “rental contract composition effect” described in section 3, then the optimal stopping threshold is $\bar{o}(r, d) = \infty$, i.e. it is *never optimal to sell an existing vehicle*. These results are due to the assumption that average daily maintenance costs EM do not increase as a function of odometer value, and that rental rates do not decrease as a function of odometer values. While there is substantial empirical justification for these assumptions *over the range of our observations*, it is questionable that these assumptions will continue to be valid as a vehicle's odometer and age increases indefinitely, far beyond the range for which we have any observations.

In order to make headway, we proceed to calculate the optimal replacement policy under *extremely conservative assumptions about increases in maintenance costs and decreases in rental rates beyond the range of the data*. That is, we will assume that beyond the range of the observations, maintenance costs increase at a very rapid rate as odometer increases, and that to induce customers to rent older vehicles, daily rental rates must be steeply discounted. Consistent with the data, over the range from $o \in [0, 130,000]$ kilometers, we assume daily maintenance costs do not increase, however outside the range of the data, we assume that daily maintenance costs start increasing at a very rapid rate, reaching a level that is *11 times* the daily maintenance costs of vehicle with 130,000 kilometers by the time the vehicle reaches 400,000 kilometers.

For rental rates, we assume that in order to induce consumers to rent older vehicles, the company must reduce the daily rental rates on the older vehicles in its fleet at a rate that is linear in the vehicle's odometer value. We assume a very steep decline in rental rates, so that at the point a vehicle reaches 400,000 kilometers the daily rental rate would be *zero*. For a vehicle with 265,000 kilometers, the rental rate it can charge is only 1/2 the rate it charges for vehicles that have 130,000 or fewer kilometers on their odometers. As we noted, the firm does in fact have a small number of vehicles in its fleet with odometer values in the range $(130000, 265000]$ yet it does not offer discounts on rentals of these vehicle and nevertheless still

succeeds in renting them to customers. We view this as evidence that the rental discount function that we have assumed is actually much steeper than necessary to induce some of the firm's customers to rent older vehicles.

Cho and Rust (2007) provide figures illustrating the calculated optimal replacement thresholds $\bar{o}(d, r, \tau)$ and the associated value functions $V(d, r, o, \tau)$ for the same three car types analyzed here. Although we do not repeat these figures here, it suffices to relay the general finding that the thresholds are roughly twice as large as the mean odometer value at which the company currently replaces its vehicles. The optimal replacement thresholds are roughly at 150,000 kilometers whereas the company currently replaces its vehicles after approximately 75,000 kilometers.

For all three car types, the "rental value", i.e. the difference in the value of keeping rental car versus selling it for a new one is a steeply decreasing function of odometer value. This is partly due to the steep decrease in the resale value of a car except that *the depreciation in a vehicle's rental value is an even steeper function of its odometer value than its resale value*. This result is an interesting contrast to the relatively mild "aging effects" that we found in the econometric analysis in section 3. Note that our assumed sharp drop off in rental rates and sharp rise in maintenance costs do not start until a rental car's odometer reaches 130,000 kilometers, yet the decline in the rental value of a car occurs immediately. The only aging effects we uncovered in the econometric analysis before 130,000 kilometers was a very mild tendency for cars to switch from long term contract to short term contracts, and for the fraction of the time they spend idle on the lot waiting to be rented gradually increases. But this "rental contract composition aging effect" is not steep enough to explain the sharp declines in the rental values of rental vehicles.

The key explanation for the rapid drop in the value of a rental car as a function of odometer is the *horizon effect*. Essentially, the instant a company purchases a new car, it represents an large investment that will be earning the company a stream of profits for a finite period of time until the car reaches its replacement threshold at which time the first will have to incur another large expenditure to buy another new vehicle. The value of keeping an existing car depends on the *expected future profits* over the life of the car, but the new purchase price of the *current* vehicle is treated as a "sunk cost." As the vehicle's odometer increases from zero towards the optimal replacement threshold, the expected discounted value of remaining profits *for the current car* necessarily decreases since the remaining life *of the current car* decreases quickly. When the company finally replaces the vehicle it must incur the cost of buying a new replacement vehicle and the process starts over again.

Note that the difference between the value of keeping a (just purchased) new vehicle and immediately

trading it for another new vehicle, represents the expected discounted profit that the firm expects to earn on the current vehicle over its lifetime. For the luxury car type, the value of keeping a newly acquired brand new car is \$375,000, whereas the value of immediately selling it is \$366,000. Thus, *the company expects to make a net discounted profit of approximately \$9,000 over the service lifetime of a single luxury vehicle.* The total discounted profits are higher, \$375,000, since this is the expectation of discounted profits earned from an infinite sequence of rental vehicles. Note that the \$9,000 figure represents *discounted profits* for a single vehicle. The mean total *undiscounted profits* over the life a compact vehicle are \$35,970, which is nearly three times larger than the actual mean undiscounted profits, \$13,719, under the *status quo*. However, as we discussed above, we think it is misleading to compare policies in terms of their impact on profits of a single vehicle, since it keeping a car longer would nearly always increase the profitability of the *current vehicle* but fail to account for the forgone higher profits that could be earned on the *next vehicle*. By comparing expected discounted profits over an infinite horizon, we can properly take the effects of extending the age at replacement on the profits of all future vehicles into account.

In order to learn more about the implications of the optimal replacement policy, particularly about the distribution of ages at which replacements occur, we compared simulated outcomes under the optimal replacement policy to simulated outcomes under the *status quo*. We found that under the optimal replacement policy, the mean odometer value at replacement is more than twice as large as the mean value under the *status quo*. The variance in odometer values about the mean value is also less under the optimal replacement policy than under the *status quo*, the mean age at replacement ranges from 4.6 to 5.0 years under the optimal replacement policy versus being between 2.6 to 2.7 years under the *status quo*.

We emphasize that it is optimal to keep these vehicles longer despite the rather substantial increases in maintenance costs and reductions in rental rates that we have assumed occurs after 130,000 kilometers. Indeed, almost all replacements that occur under the optimal replacement policy occur well after 130,000 kilometers, when these “adverse” aging effects have kicked in. Note, however, that all of the cars are replaced before they reach 265,000 kilometers, which is the point where rental rates are discounted to 50% of the rate for a vehicle with 130,000 kilometers. Also, according to our assumptions daily maintenance costs are about 5 times higher for vehicles at 265,000 kilometers than the values for vehicles that have fewer than 130,000 kilometers. So the combination of the rental discounts and rapid increase in maintenance cost greatly alter the optimal replacement policy. Instead of being optimal to *never* replace its existing vehicles, the assumptions of rapidly rising maintenance costs and rapidly declining rental rates after 130,000 kilometers lead to a finite optimal replacement threshold. It is surprising is that despite these

extremely pessimistic assumptions, the optimal replacement policy still entails keeping cars about twice as long as the company currently does.

In order to compare the discounted profits under an infinite horizon, we need to make extrapolations of the firm's *status quo* replacement policy into the indefinite future. Since the data obviously only covers a relatively short time span of the firm's operations, it is clear that certain "forecasting assumptions" are required. In particular, we assume *stationarity*, i.e. that real prices of cars, maintenance costs, rental rates and so forth will be constant for the foreseeable future. Under these (strong) assumptions, we can calculate the value of an *infinite sequence of replacement vehicles over an infinite horizon*.

To simplify the analysis, we focus on comparing the value of a newly purchased brand new car that has just entered the lot. In Table 4 we report the value of a new car that has just arrived in the lot under the optimal replacement policy, and compare it the value of a new car that has just arrived in the lot under the the firm's *status quo* operating strategy. Table 4 also presents an "equivalent daily profit rate" which is approximated as $(1 - \beta)V(0, r_0, 0)$ and $(1 - \beta)V_\mu(0, r_0, 0)$ where $\beta = \exp\{-r/365\}$ is the daily discount factor. For our calculations we have assumed that $r = .03$, and this implies a daily discount factor that is quite close to 1, $\beta = 0.99991781$. Our conclusions are robust to fairly wide changes in the discount rate.⁷

The first section of Table 4 presents the expected discounted values and the daily expected profit equivalent values for the optimal replacement policy for each of the three car types that we analyzed. Also, to provide an measuring stick for these numbers, the top line also presents the average price of a new vehicle for each car type. We see that for the compact car, for example, the expected present discounted value of profits is \$268,963, which is 27.8 times the cost of a new compact car. Applying the final value theorem, we find that this discounted profit is equivalent to about \$22.11 in profits on a daily basis. The second section of Table 4 presents the expected discounted value of profits under the *status quo*. The expected discounted value of profits over an infinite horizon is \$196,589, which is equivalent to \$16.16 on a daily basis. Thus, we see that according to the model's predictions, the firm could increase its discounted profits by 38% (i.e. $V(0, r_0, 0)/V_\mu(0, r_0, 0) = 1.37$), if it adopted the optimal replacement policy, in combination with "deep discounts" in rental prices of older vehicles.

We find that for the luxury car type, the firm's replacement strategy is closer to optimality: its profits would increase by 18% under the optimal replacement strategy. However for the RV, the firm's existing

⁷ According to the *final value theorem*, (see, e.g. Howard, 1971, p. 46) for any convergent sequence $\{a_t\}$ we have $\lim_{\beta \rightarrow 1} (1 - \beta) \sum_{t=0}^{\infty} \beta^t a_t = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T a_t$. There are stochastic extensions of this result that imply that for β close to 1, $(1 - \beta)V(d, r, o)$ is close to the "long run average profits", which in our case corresponds to an equivalent daily profit. The statement of the stochastic version of the final value theorem is more complex and is omitted here, but the basic result is the same as the deterministic version of the final value theorem given above.

policy appears to be far from optimal: the present discounted profits are predicted to be *2.4 times higher under the optimal replacement strategy*.

	Compact All Locations	Luxury All Locations	RV All Locations
Quantity	Value	Value	Value
\bar{P}	9668	23389	18774
Expected Discounted Values Under Optimal Replacement Policy			
$V(0, 0, r_0)$	268963	374913	327057
$(1 - \beta)V(0, 0, r_0)$	22.11	30.81	26.88
$V(0, 0, r_0)/\bar{P}$	27.8	16.0	17.4
Expected Discounted Values Under <i>Status Quo</i> Replacement Policy			
$V_\mu(0, 0, r_0)$	196589	318247	136792
$(1 - \beta)V_\mu(0, 0, r_0)$	16.16	26.16	11.24
$V(0, 0, r_0)_\mu/\bar{P}$	20.3	13.6	7.3
Ratio of Expected Values: Optimal Policy versus <i>Status Quo</i>			
$V(0, 0, r_0)/V_\mu(0, 0, r_0)$	1.37	1.18	2.39

Table 4: Comparison of Profits/Returns: Optimal Policy versus *Status Quo*

We undertook another set of discounted profit calculations to see if our conclusions are robust to even more conservative assumptions about the increase in maintenance costs and required discounts in rental rates. Under this even more conservative scenario, we assumed that maintenance costs begin rising steeply even earlier, at 60,000 kilometers. We also assumed that rental rates would have to start declining after 60,000 kilometers at even a faster rate than we previously assumed, so that by the time a car reaches 210,000 kilometers, its daily rental rate would be *zero*. As expected, it is optimal to replace cars even sooner under this more pessimistic scenario. Nevertheless, the optimal replacement policy still entails keeping cars roughly twice as long as the company currently keeps them, and even under these extremely pessimistic assumptions. Note that our pessimistic assumptions do not apply in our calculation of discounted profits under the *status quo*, nevertheless the optimal replacement policy *still results in significantly higher profits than the status quo*. We see that expected discounted profits for the Compact, Luxury and RV vehicle types are predicted to increase by 25%, 6% and 100%, respectively. Thus, our predictions are quite robust to variations in the assumptions. Our estimates of the increases in profitability from delaying replacements of rental vehicles are likely to be fairly conservative: most likely the company would not need to discount rental rates as steeply as we have assumed, and if so, the gains it would realize from adopting an optimal replacement strategy would be even larger than we have estimated.

6 Conclusion

We view this paper as providing primarily a *methodological* contribution, but as we discuss below, it is still not clear whether the econometric model we developed will have *practical* value. The methodological contribution is to show how to integrate econometric duration models and (regenerative) optimal stopping theory in order to evaluate the profitability of the operating strategy of a firm. We have shown how this apparatus can be used to test the hypothesis that the firm is profit maximizing and we have provided evidence that the rental car company that we analyzed here is not maximizing discounted profits even though we did show that the company is highly profitable.

Of course we always have to keep in mind the possibility that the *model is wrong* and not interpret any deviation between the model predictions and actual outcomes as *prima facie* evidence that the firm is not behaving optimally. Instead, the discrepancy should be treated as *prima facie* evidence that there is something wrong with our model. Fortunately, there is a further way to test this: namely by undertaking an experiment to test whether the predictions of the econometric model are correct. The rental car company was sufficiently convinced by the analysis in this paper to undertake a small scale experiment to test the main predictions that keeping rental cars longer would increase the firm's profits. We view this as a limited practical success.

We do not have sufficient space here to describe this experiment in detail, but the experiment involved a combination of discounting the rental prices of older cars and keeping the cars longer before replacing them. We refer the reader to Cho and Rust (2008) for an analysis and interpretation of the outcome of this experiment.

At this point, this paper should be regarded primarily as a concrete illustration of the role that econometrics *might* be able to play as a decision tool that could assist companies in improving outcomes (e.g. profits) through the use of structural models as a tool for simulating and systematically searching for improved policies via dynamic programming methods and computer simulations (i.e. computational experiments) rather than relying on *ad hoc* trial and error methods, (i.e. real experiments) that are very costly and time consuming. We still believe that a combination of econometric modeling, computational experimentation, and real world experimentation can be extremely effective: the econometric modeling and computational experiments can be used to search for attractive policies at very low cost, but real world experiments should be undertaken to verify that predictions of the econometric model are accurate.

We feel that the structural approach taken in this paper, i.e. developing an econometric model that

can simulate the operations of the company both under the *status quo* and under a range of counterfactual alternative replacement policies, is both more ambitious and promising than the reduced form econometric methodology that currently dominates applied econometric work. A severe limitation of the reduced-form methodology, including the approach taken in the “treatment effects” literature, is that it is fundamentally backward looking. The goal of most of these studies is only to try to estimate the “treatment effect,” typically treated as unidimensional variable, for some policy change that has been taken in the *past*. The treatment effects literature interprets significant historical changes in policies as quasi-random “policy experiments” that can serve as nearly exogenous policy shifters or “instruments” for determining “causal effects” in their econometric analyses. However, it does not attempt to predict how new, hypothetical policy changes might affect outcomes in the *future*.

Most of the practitioners in the reduced-form and treatment effects literatures are unwilling to undertake the econometric modeling necessary to predict the outcomes of a range of hypothetical counterfactual policy experiments in order to give policy makers guidance on the policy experiments of most interest (e.g. policy changes that optimize some well defined criterion). As a result, the reduced-form and treatment effects approaches are not very useful for practical decision making, at least in a real time, forward looking environment. We believe that increasingly policy makers need assistance to help them make better decisions going forward, and that it is not enough to simply look backward and evaluate the outcomes of decisions they took in the past, although we certainly don’t deny that there is much to be learned from evaluations of the success or failure of past decisions.

We believe that while structural econometric modeling is still in its infancy, the models are rapidly evolving and improving and may soon be sufficiently realistic and accurate to be able offer forward looking guidance to policy makers to actually help them make better decisions. But clearly the path to success will be littered with many failed attempts along the way. While it is still too early to determine whether our model and policy forecasts will be a practical success in this case, we would like to acknowledge the promising results obtained from previous studies by Paarsch and Shearer (1999) and Shearer (2004) who used structural economic models to design alternative compensation schemes for firms and designed experiments to verify the predictions of their models, and Todd and Wolpin (2006) who demonstrated that structural econometric models can provide accurate out of sample predictions of the treatment effects generated by randomized experiments.

References

- [1] Cho, Sungjin (2008) “An Empirical Model of Mainframe Computer Investment” *Journal of Applied Econometrics* forthcoming.
- [2] Cho, Sungjin and J. Rust (2008) “The Flat Rental Puzzle” manuscript, University of Maryland.
- [3] Howard, R. A. (1971) *Dynamic Probabilistic Systems: Volume II: Semi-Markov and Decision Processes* John Wiley & Sons, Inc. New York.
- [4] Levitt, S. (2006) “An Economist Sells Bagels: A Case Study in Profit Maximization” manuscript, University of Chicago.
- [5] Paarsch, H. and B. Shearer (1999) “The Response of Worker Effort to Piece Rates” *Journal of Human Resources* **34-4** 643–667.
- [6] Rothwell, G. and J. Rust (1997) “On the Optimal Lifetime of Nuclear Power Plants” *Journal of Business and Economic Statistics*. **15-2**, 195–208.
- [7] Rust, John (1985) “Stationary Equilibrium in a Market for Durable Assets” *Econometrica* **53-4**, 783-805.
- [8] Rust, John (1987) “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher” *Econometrica* **55-5**, 999-1033.
- [9] Rust, John (1994) “Structural Estimation of Markov Decision Processes” in R. Engle and D. McFadden (eds.) *Handbook of Econometrics* Volume 4, 3082–3139, North Holland.
- [10] Rust, John (1996) “Numerical Dynamic Programming in Economics” in H. Amman, D. Kendrick and J. Rust (eds.) *Handbook of Computational Economics* Elsevier, North Holland.
- [11] Rust, John (2007) “Is Econometrics Useful for Public Policy Making? A Case Study of Disability Policy at the Social Security Administration” manuscript, University of Maryland.
- [12] Shearer, B. (2004) “Piece Rates, Fixed Wages and Incentives: Evidence from a Field Experiment” *The Review of Economic Studies* **71-2** 513–534.
- [13] Todd, P. and K. Wolpin (2006) “*Ex Ante* Evaluation of Social Programs” Penn Institute for Economic Research working paper 06-022.