When to Sell the Car? Actual Replacement Decisions of a Rental Car Company versus Predictions of an Optimal Stopping Model

Sungjin Cho, *Hanyang University*
John Rust, *University of Maryland**

April 4, 2006

Abstract: We study a large, successful rental car company that buys new vehicles and rents them to a series of customers who choose long or short term rental contracts. The company earns extraordinarily high rates of return on its rental cars, with average internal rates of return between purchase and sale of approximately 50%. We develop a stochastic semi-Markov duration model of the company's operations, and via stochastic simulations, show that the model closely mimics the outcomes experienced by this company under its current operating strategy, which is to charge "flat" rental rates (i.e. rates that do not decline with vehicle age or odometer) and to replace old cars with new vehicles after approximately three years of operation. In particular, our model replicates the high internal rates of return earned by this company on its rental vehicles. Using this model, we can evaluate the profitability of a wide range of alternative replacement and rental pricing strategies. We characterize the profit maximizing replacement strategy by formulating the problem of periodic replacement of vehicles in the company's fleet as a *regenerative optimal stopping problem*, and solve the problem numerically. Depending on the vehicle type, we find that the company's expected discounted profits would be between 6% to over 140% higher under an alternative operating strategy where vehicles are kept longer and rental rates of older vehicles are discounted to induce customers to rent them.

Keywords: automobile rental/leasing, duration models, semi Markov processes, 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 and Finance, Hanyang University, Seoul, Korea phone: 82-2-2220-1019, email: `sungjcho@hanyang.ac.kr`. The authors are grateful to the National Science Foundation and the Pittsburgh Supercomputer Institute `www.psc.edu` for computing resources that helped make this work possible. This paper is preliminary and should not be quoted without the permission of the authors.

1 Introduction

This study presents a detailed analysis of the replacement decisions of a large, successful rental car company that buys new vehicles and rents them to series of customers who choose between long and short term rental contracts. Besides the size and composition of its rental fleet and the rental rates it charges, a key operating decision that the company makes is the timing of vehicle sales. We focus on the latter decision, and the costs and benefits of keeping a rental car longer versus the alternative of selling it relatively quickly. Due to the rapid depreciation in the resale prices of cars, replacing vehicles too quickly can reduce profits by increasing trading costs and reducing rental revenues.

This company is by any measure extremely successful in the rental car business: the average gross internal rate of return it earns on cash flows from the initial purchase of a vehicle to its ultimate sale is approximately 50%.[1]

In economic theory, the “standard hypothesis” is that firms choose operating strategies that maximize the expected discounted value of cash flows arising from their operations. Given how successful this firm appears to be in the rental car business, it is natural to suppose that this firm is maximizing profits. This is the basic hypothesis we intend to “test” in this analysis. To do this, we develop a stochastic semi-Markov duration model of the company’s operations, and via stochastic simulations, show that the model closely mimics the outcomes experienced by this company under its current operating strategy, which is to adopt “flat” rental rates (i.e. fixed per day rental charges that do not decline with the age or odometer value of the vehicle) and to replace old vehicles with brand new ones after approximately three years of operation. In particular, our model replicates the high internal rates of return earned by this company on its rental vehicles. Thus, we argue that our stochastic model provides a good representation of the operations and replacement decisions of this company under its *status quo* operating strategy.

Using our model, we can evaluate the profitability of a wide range of alternative replacement and rental pricing strategies. Conditional on certain assumptions about rental rates (i.e. namely that they are fixed at the firm’s current values, except for discounts offered to customers who choose to rent older vehicles) and replacement decisions (i.e. we assume that when the firm replaces an existing vehicle it purchases a

[1] The internal rate of return is defined as the interest rate that equates the present value of cash inflows and cash outflows — i.e. it is the “breakeven” discount rate where the net present value of an investment in a rental car equals zero. The reported internal rates of return are based on gross cash flows, i.e. they have not be adjusted to account for taxes and administrative and insurance costs. Internal rates of return based on net after tax cash flows would be somewhat lower than the gross internal rates of return, but are still extremely high. Our impression is that these rates of return are much higher than typical rates of return earned on investments involving comparable degrees of risk.

brand new replacement vehicle, rather than a slightly used vehicle at a large discount over the price of a brand new car), we characterize the profit maximizing replacement strategy by formulating the problem of periodic replacement of vehicles in the company's fleet as a *regenerative optimal stopping problem*. We employ numerical methods to compute the optimal replacement strategy, using the econometrically estimated resale price functions, maintenance costs, and hazard rates and transition probabilities governing movement of cars between long and short term rental spells and lot spells. It turns out that the optimal strategy predicted by our model is sensitive to assumptions about *aging effects* of vehicles — effects that are hard to extrapolate outside of our current sample of data due to the fact that this company replaces its vehicles relatively frequently, typically after three years of service life.

The most obvious aging effect is the rapid decline in the resale price of a car as a function of its age (or odometer value). However with respect to other measures of the condition of the car such as maintenance costs or the duration of rental spells, our econometric analysis finds few other clear aging effects for vehicles in this company's fleet, at least within the three year horizon over which it keeps them. Thus, although customers prefer newer cars to older ones, the company does not discount daily rental rates as a function of the age or odometer of the vehicle so there there is no "aging effect" in its rental rates. We also do not find any evidence of an upward trend in maintenance costs as a function of the age and odometer of the vehicle. Lastly, we do not find any evidence of age or odometer dependence in mean duration of vehicles in rental contracts, or for cars in lot spells waiting to be rented.

Beside the rapid depreciation in resale price, the only aging effect that we find from our econometric analysis is that newer cars tend to start out in long term rental contracts at the start of their life and are gradually switched to an increasing share of short term rental contracts toward the end of their lives. We find that long term contracts have a high probability of "roll over" and thus, the cars which start out in long term rentals tend to spend less time in the lot compared to cars that are in short term rentals. Thus, as vehicles age, there is an increasing probability that they will be rented in short term contracts, and as a result, there is a higher probability that older car will be idle due to the greater probability of being in a lot spell between successive short term rental spells.

However if this "contract composition age effect" and the rapid depreciation in resale values are the only aging effects, our optimal stopping model predicts that the company should *never* replace its vehicles. Instead, the optimal replacement policy entails maintaining the existing stock of vehicles indefinitely. However this result depends on the assumption that consumers do not mind driving arbitrarily old vehicles, and that the company does not incur rapidly increasing maintainance costs to enable them to "live" forever.

Company executives believe that its customers expect to rent "new" vehicles, and this appears to be the main consideration underlying their three year vehicle replacement target. However it seems plausible that many customers might not consider cars that are more than one year old to be "new" yet more 90% of the cars we studied were two years or older, and approximately 40% were three years or older. Nevertheless, the company is able to rent these older vehicles without providing any discount to its customers. While we have no direct evidence of the magnitude of the reductions in rental rates that would be required to make customers indifferent between renting "new" vehicles at the company's current rental rates and vehicles that they consider to be "old" we show that even when extremely deep (and we think implausibly large) discounts in rental rates are used to induce customers to rent older vehicles, the optimal replacement strategy entails keeping vehicles roughly twice as long (measured either in terms of vehicle age or odometer at time of sale) than the company currently keeps them.

Depending on the type of vehicle analyzed, we find that expected discounted profits would be at least 6% higher, and as much as 140% higher under this alternative operating strategy than the company's existing operating strategy. Further, the actual increase in profits could be significantly higher than the values we predict due to our extremely conservative assumptions about the size of discounts and magnitude of maintenance costs for older vehicles. We have made strong, and we think implausibly pessimistic, assumptions about the required discounts in rental rates and the rate of increase in maintenance costs in order to create a strong implicit bias against holding older vehicles. Nevertheless, the model predicts that under an optimal strategy, the firm should keep its vehicles for much longer than it currently does.

Specifically, we have assumed that maintenance costs suddenly increase at a very rapid rate after 130,000 kilometers, so that the average daily maintenance costs for a vehicle with 400,000 kilometers is *11 times* the average daily maintenance cost of a vehicle at 130,000 kilometers. We also assume that the firm must discount the rental rates it charges to customers to induce them to rent older vehicles. We assume that daily rental rates for both long and short term rentals are reduced at a linear rate as a function the vehicle's odometer starting at 130,000 kilometers, until they reach 0 for vehicles with 400,000 or more kilometers on their odometers. This implies that the daily rental rate for a vehicle with 265,000 kilometers *is only 1/2 the daily rental rate that the company earns for its vehicles under its existing "flat rental schedule"*. We note that company does have a small number number of vehicles with more than 265,000 kilometers on their odometers, but nevertheless rents them at the full daily rental rates. This, combined with the evidence provided by the company's "experiment" suggests that it is unnecessary to reduce rental rates by this much to give customers sufficient incentive to rent older vehicles.

Indeed, the company did previously experiment with discounts in rental rates for older vehicles, but discontinued the experiment because too many customers were choosing older vehicles in preference to renting the newer cars in their fleet. The experiment involved a 20% discount in daily rental rates for any vehicle over two years old. We view this as evidence that 20% discount was too large, and that with lower discounts, the company could succeed in inducing customers to rent old vehicles without jeopardizing its ability to rent its newer cars at the existing full rental rates. We believe that a moderate age or odometer-based discounting strategy could lead to a "win-win" situation: it could enable the company to increase its profits while at the same time providing a wider range of choices and benefits to its customers. Customers would benefit since they can always choose the default of renting a newer car at the full rental rate. However many customers may prefer to rent an older vehicle at a reduced rental rate, and these customers will be strictly better off under this alternative rental rate structure. The company would benefit from being able to keep vehicles in its fleet longer, and thus earn more rental revenue over a longer holding period that would help to "amortize" the high trading costs it incurs from the rapid initial depreciation in vehicle prices. Thus, by appropriately discounting its rental rates, the company should be able to significantly increase its profits without risking its reputation and the good will of its customers.

Section 2 describes the rental car data. Section 3 presents an econometric model of the company's operations and provides our econometric estimates of the our models of vehicle resale prices, durations in rental and lot spells, and transitions between spells. Section 4 shows, via stochastic simulations, that this model provides a good approximation to the company's behavior/outcomes under its *status quo* operating strategy. Section 5 formulates and solves a dynamic programming problem that provides the profit-maximizing operating strategy under certain "maintained" assumptions about rental rates, and other constraints on replacement decisions by this company. Section 6 compares actual and predicted optimal operating strategies, discounted profits, and rates of return for 3 different makes, models and rental locations of vehicles. Section 7 discusses extensions of this analysis to issues of optimal selling strategy (i.e. whether the company should sell one of its vehicle regardless of what price is offered for it, or whether the company should adopt a "reservation price" sales strategy to achieve a higher resale value for its rental cars), and fleet allocation (i.e. whether the company should reallocate its portfolio of rental cars, increasing holdings of certain types of cars and decreasing holdings of others). Section 8 presents conclusions, including our views of the contributions and limitations of this analysis and some of the issues that arise in the experimental and quasi-experimental validation of the predicted increase in profits from our suggested changes in the firm's operating strategy.

2 Data

We obtained data from a large syndicated rental car company that owns and rents a large fleet of rental vehicles in over 100 different locations in its region of operation, which include urban locations (e.g. car rental facilities in main cities and near major airports) and tourist areas. Due to confidentiality agreements we have with this company, we are unable to disclose the name of the company, the rental syndicate it is part of (e.g. Avis, Hertz, National/Alamo, etc.) or further details about its location.

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. Purchase times of individual vehicles are fairly evenly spread out in time. 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 consist of information on date and purchase price for each vehicle it 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, where rental agents appear to have filled in rough estimates of the out and in odometer values at the rollover dates where the customer decided to keep the 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 travelled during each rental contract we relied on some (we believe reasonable) econometric modelling assumptions that we will describe shortly.

The company also provided us with 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 the 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 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. The two most important predictor variables (besides the make/model of vehicle) are the vehicle's age and odometer value at time of sale.

The firm 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 *roll over* (i.e. renew) the previous contract for another month. Rental contracts are at a daily rate with no additional charges for distance travelled 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.

We are not familiar with the exact terms of the long term rental agreement, but we presume that the company retains the right not to automatically roll over a contract at the end of a previous one if it decides to sell the vehicle. Thus, in our model, we assume that on the first day of any long term contract, the company can decide to sell the car and provide the customer with a substitute vehicle of the same make and model. However if the company decides not to sell the car on the date a contract rolls over, it must wait until the end of the next contract (or when the car comes back, if the customer decides to return it before the end of the contract period) before it can sell the car. We assume that the company can sell the car at any time it pleases if it is in a lot spell, i.e. not currently in the middle of short term or long term rental spell.

Figure 1 illustrates typical rental histories for three different cars in the company's fleet: 1) a compact

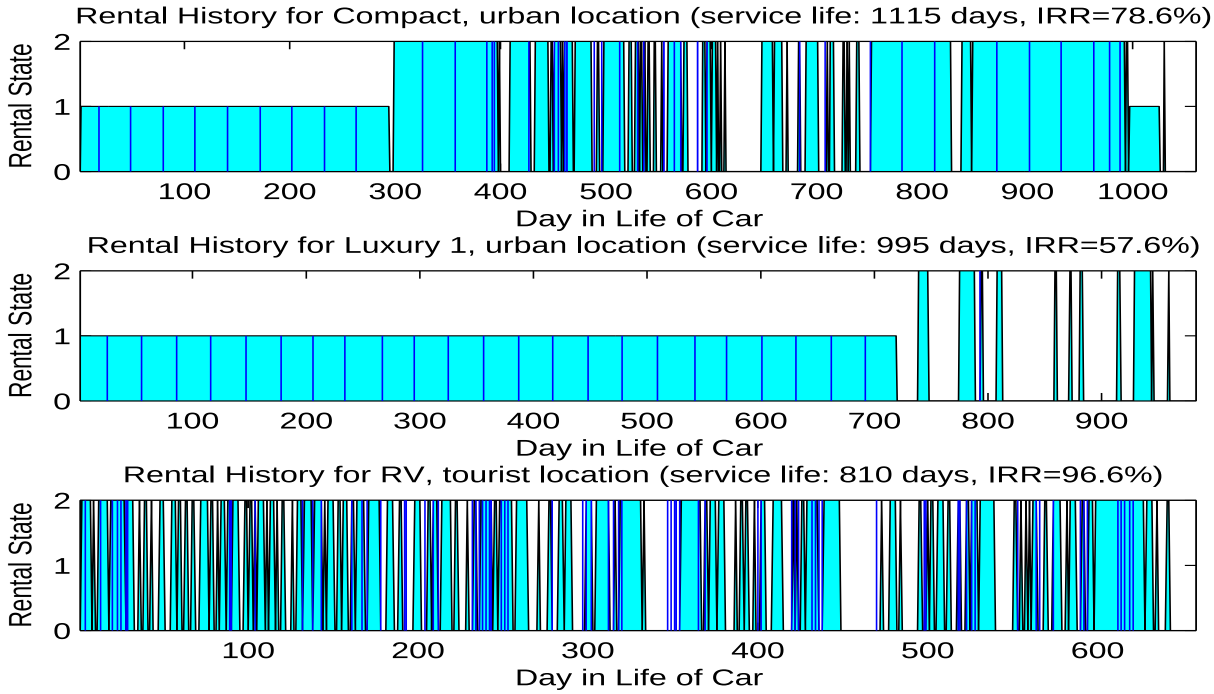

Figure 1 Typical Rental Histories for Three Cars in the Company Fleet

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. Quite possibly the succession of unbroken long term contracts could represent the same customer who rolled over their monthly contracts into the *de facto* equivalent of a lease, lasting nearly one year in the case of a 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:

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

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 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. Thus, the compact car experienced relatively greater price depreciation, but this could also be due to its being 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).

Figure 2 shows the distributions of internal rates of return for the same three classes of vehicles for which individual rental histories were plotted in Figure 1. These distributions are for all cars for which we have complete rental histories and include cars from all rental locations.[2] We see that the high IRRs shown in the particular rental histories in figure 1 are not atypical, although we do see a significant level of variability in realized rates of return, and differences across the three car classes in the return distributions.

It is apparent in figure 2 that the compact class has the highest average rates of return, with a median

[2]When we restrict the samples to particular rental locations, such as large urban locations, we get fewer number of observations, but the sample restriction does not have a signifi cant effect on the distributions.

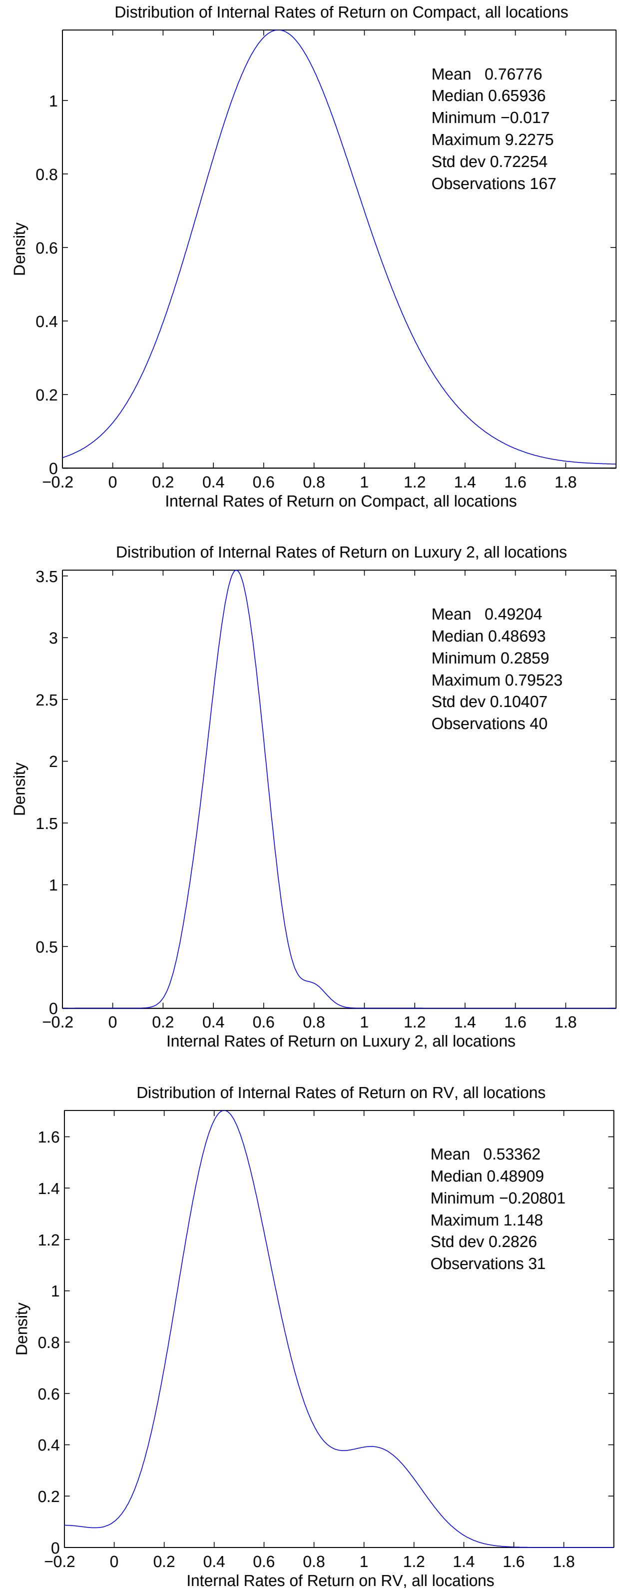

Figure 2 Distribution of Internal Rates of Return: Compact, Luxury, and RV – All Locations

IRR of 66%. The luxury and RV classes have median IRR's of 49%. There is more variance in the distribution of IRRs for the compact class, and the distribution of IRRs for RVs reflects skewness caused by a few large positive IRRs (of 300% and 900%, respectively). Clearly, rentals have a stochastic (i.e. unpredictable) component, and certain cars could achieve high rates of return if they happened to be "in the right place at the right time." That is, a car that was rented for most of its lifespan would be expected to have a higher IRR than a car that experienced considerable idle time on the lot. However variations in new purchase prices, proceeds from sales of the car, and variability in total maintenance costs are also other obvious stochastic factors that can affect profits and realized rates of return.

Besides the net loss from price depreciation (i.e. the difference between the price of a new car less its resale value at time of sale), perhaps the most important factor predictive of the profitability of a rental car is its *effective capacity utilization.* We define this as the fraction of the car's service life that it was rented. In figure 1, the luxury car had the highest capacity utilization rate, having been rented 775 days out of its 995 day service life, corresponding to a 78% utilization rate. The compact car had a utilization rate of 74% and the RV had a utilization rate of 42%. It might be expected that a recreational vehicle in a tourist location would have a lower utilization rate, reflecting idleness in off-peak seasons and weekdays and non-holiday times.

What accounts for the high rate of return on the RV even though its capacity utilization rate was lower? The average daily rental rate for the RV was $114, compared to only $34 per day for the compact, and $81 per day for the luxury car. Also, these rental rates are for *short term* contracts. The daily rental rate earned on long term rentals is even lower: $20 per day for the compact and $44 per day for the luxury car. Thus, it is not only capacity utilization that matters, but also the fraction of time spent in short term rental contracts. Short term rental contracts are more lucrative in terms of the higher daily rental rates, but their average duration is much shorter, and there is more idle time associated with these contracts due to the higher probability a car will be on the lot between successive short term rentals. Also, the higher turnover rates of short term rentals lead to higher costs of cleaning/maintaining the car after each short term rental spell in order to make the car ready for the next rental.

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 on rental vehicles. We report three regressions for the vehicle types: compact, luxury, and recreational vehicle, pooling over all rental locations.[3] The predicted signs of the coefficients are mostly consistent with intuition: the utilization rate

[3]The results are basically unchanged if we use the logarithm of the internal rate of return as the dependent variable:

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.

	Compact All Locations		Luxury All Locations		RV All Locations	
Variable	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

There are three variables for which the signs of the regression coefficients are *a priori* ambiguous: the fraction of the rented life in long term contracts, and the age and odometer value on the vehicle at time of sale. The reason why the first of these is ambiguous has been discussed above: even though daily rental rates for short term contracts are significantly higher than for long term contracts, the mean duration of short term rental spells is lower and the likelihood of idle time on the lot between successive short term rentals is higher. Thus it is not clear *a priori* whether the firm would prefer its rental cars to spend a larger fraction of their rented life in short term or long term contracts.

For most the variables where we do have unambiguous expectations about how they affect the IRR, the results generally confirm our expectations: higher purchase prices reduce the IRR, higher resale increase it, and the coefficient on the capacity utilization rate is also positive and statistically significant. The signs on the daily rental rates are also positive and generally statistically significant. The fraction of the time the car was rented long term has a negative coefficient, suggesting long term contracts are *less* profitable than short term contracts on a per day basis, and the utilization rate has a positive estimated coefficent (as expected) which is significant for the luxury and RV car types, but insignificant for the compact.

the R^2 statistics are slightly lower but the same pattern of signs and signifi cance levels for the coeffi cients emerges for this alternative specifi cation for the dependent variable in the regression.

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 coefficients of age, odometer, and maintenance costs have variable signs and are frequently statistically insignificant. One reasons is that these variables have a high degree of collinearity, especially age and odometer. However 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. In this case age has a positive coefficient and total maintenance cost has a negative coefficient. But even in this case, the effect of age on IRR is small: the regression results predict that keeping a luxury car for 100 more days increases the IRR by 0.03.

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. If the company has chosen the optimal odometer threshold o^* at which to sell the car, then at the optimal threshold o^* we have

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

It follows that if the odometer values at which the company sells its cars are approximately eqal 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^* at time of sale on the realized 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.[4]

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 our 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, we are not aware of any recalls or any factors or variables that we could exploit to use an instrumental variables approach.

Thus, there are 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 a specific change in the company's replacment policy, either replacing cars earlier or later than they do under the *status quo,* and the remaining cars of the same car type at the same location would be assigned to the *control group* and would continue to be subject to the company's existing or *status quo* operating 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

[4]The 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.

(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 treatment could involve replacing cars earlier or later that under the *status quo* and there could be many possible choices for how much earlier or later, and whether the appropriate threshold should be based on vehicle age or odometer value. One needs 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 the 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 our 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 of the Rental Car Company

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, where the previous rental state was a long term rental spell,

4. In the lot waiting to be rented, where the 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 case, we formulate the problem in discrete time, with the relevant time unit being a day. We 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 $\{1,2,3,4\}$.[5]

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 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 vehicles *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 our empirical analysis below, it turns out that a vehicle’s age t is strongly correlated 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 try to model, and since excluding the age variable results in a significant reduction in the dimensionality of the problem, we have opted to exclude

[5] Actually we could distiguish 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.

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 as we will see shortly, the model we construct does in fact keep track of vehicle age, and simulations of our model are able to accurately predict the distribution of ages at which the company sells vehicles, even though we restrict the company's decision rule about selling a vehicle to be a function of 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 realized variables 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 key components

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,
6. A model of the company's *selling decision*, i.e. the factors that motivate it to sell a given car at a particular point in time.

We will now discuss each of these components in turn, describing the econometric model we chose and the result from estimating it. The first model is a model of a car's resale price. We have data on both the new price $\overline{P}(\tau)$ 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.* Our econometric analysis will focus on three particular car types: 1) a compact, 2) a luxury, and 3) an 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. But we wish to emphasize that *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.*

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

$$\log(\overline{P}(\tau)/P_t(o_t,\tau)) = \alpha_1(\tau) + \alpha_2(\tau)o_t + \varepsilon_t. \tag{3}$$

The results from this model be interpreted as a regression with cartype-specific "depreciation coefficients" $(\alpha_1(\tau),\alpha_2(\tau))$ where α_2 measures the effect of odometer on the selling price of the vehicle. 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 2 below.

	Compact All Locations		Luxury All Locations		RV All Locations	
Variable	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 $\log(\overline{P}(\tau)/P_t(o_t,\tau))$

The regression results show that both age and odometer value are significant predictors of the resale price of used cars, however the incremental predictive power of adding age in addition to odometer value is not huge. Thus, for the case of the compact car, the R^2 statistic (which measures the fraction of the variance in used car prices explained by the regression) drops only slightly, from 38.9% to 37.5%, when the age regressor is removed, and for the RV, the R^2 drops from 42% to 41%. Nevertheless, we cannot reject the hypothesis that age is a significant predictor of used car prices, as we can see from the size of t-ratios of the age coefficient for all three vehicle types (where a value of plus or minus 1.96 corresponds to an approximate 5% significance level). In terms of the "economic significance" of the relative magnitudes of price depreciation predicted by the estimated age and odometer coefficients, note that for the compact car, the average number of kilometers travelled per day (including accounting for "dead time" on the lot) was 77. Thus, after 100 days, the compact would be predicted to have an odometer values that was about 7700 kilometers higher. From the results in table 1, we see that the regression model predicts the price depreciation from the extra kilometers would be $-1/2$%, whereas in terms of the predicted effect of age along, the extra 100 days would be predicted to lower the car's value by -1%. Thus, in some sense the effect of age on price seems to be about twice as large as the effect of odometer, *ceteris paribus.* Of course, both age and odometer tend to increase together, and so the combined effect of an additional 100 days and 7700 kilometers is roughly the sum of the age and odometer coefficients, or roughly a -1.5% depreciation in the value of the vehicle.

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 (62% $= \exp(-.48)$), than for the luxury vehicle (52%) or the the RV (43%). Figure 3 provides scatter plots of the resale prices for the three cars, graphed against the vehicle's odometer value at time of sale. The rapid early depreciation in the 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 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, but *ad hoc* extrapolation of what we think a very new used car (i.e. one with less than 20,000 kilometers) would sell for, instead of using the estimated regression intercepts, which we feel would greatly underestimate the resale value of a very new used car. We assumed that the "instantaneous depreciation" for a brand new car would be only 10% and then used a straightline interpolation from this value to the resale values implied by our regressions at an odometer value of 20,000 kilometers. As we can see, even at this relatively low odometer level we do have actual observations of sales and we see that the regression does accurately predict the mean resale price over the range of our data.

In our subsequent analysis we will show that our predictions of the optimal replacement policy are not sensitive to our assumptions about the precise shape of the depreciation curve for cars with odometer values of less than 20,000 kilometers. The simple reason for this is that under none of the scenarios that we analyzed would it ever be profitable for the company to sell its cars before 20,000 kilometers. Thus, whether the initial early depreciation in a car's value happens "instantaneously" (i.e. dropping from 100% to between 62% to 43% of its value the second it drives off the lot) or slightly more gradually as our linear interpolations in figure 3 suggest, is really immaterial from the standpoint of the company. This very rapid early depreciation in vehicle is evidently a cost that the company must deal with, and its method for

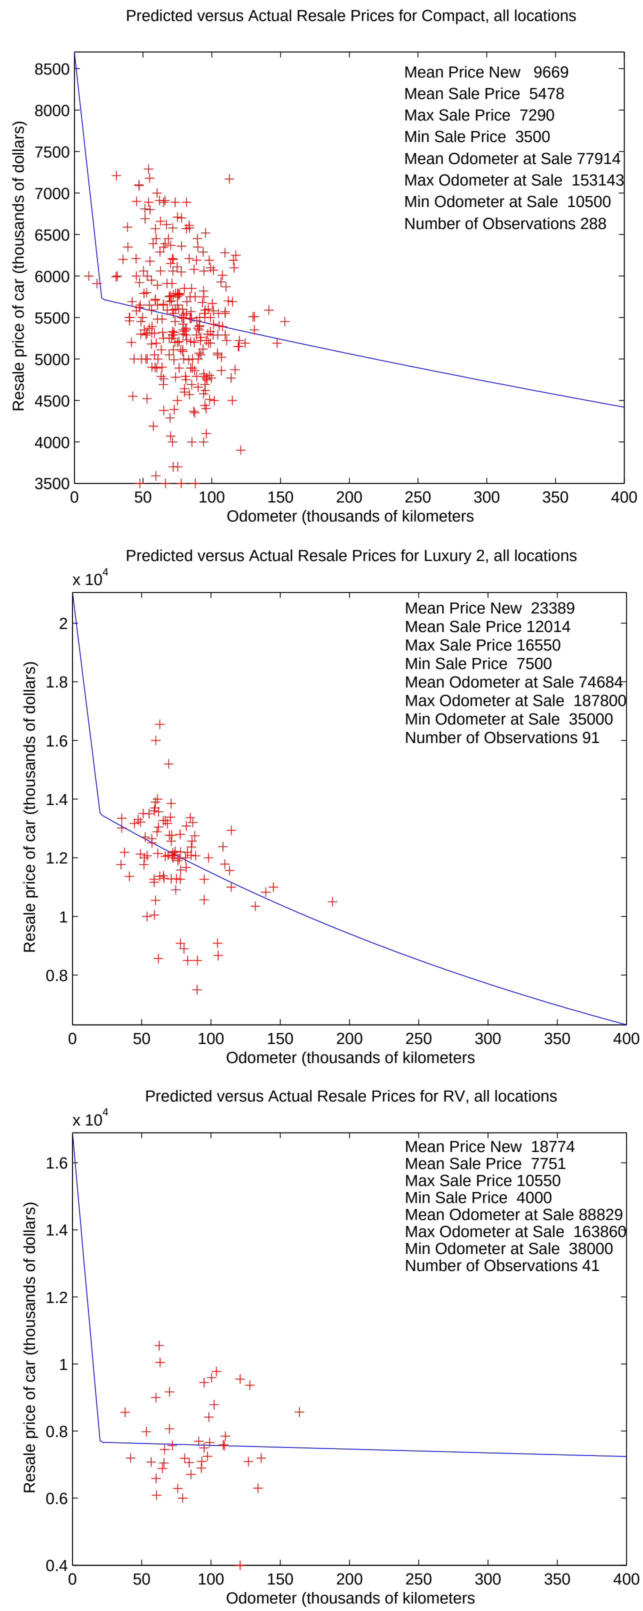

Figure 3 Predicted versus Actual Resale Prices: Compact, Luxury, and RV – All Locations

“recouping” these depreciation costs is to hold a vehicle long enough so that the rental revenue it earns far outweighs the loss the company incurs from the depreciation in resale value.

The regressions show that accidents, measured either by the number of accidents or the total cost of repairing the accidents, does not have a significant derogatory effect on the resale value of the vehicle. This is likely because of the impact of insurance, which is supposed to result in repairs following an accident that restores the vehicle to its pre-accident condition. For this reason, it is perhaps not surprising that the accident variables are insignificant.

It is perhaps more surprising that average daily maintenance cost of the vehicle over its lifetime is not a significant predictor of the resale value of a vehicle. To our knowledge, while the company is required to disclose the number of accidents that a vehicle had at the time of sale, it is not required to disclose the total maintenance costs. Thus, a potential purchaser may not have the information that a certain vehicle was a “lemon” and encountered very high maintenance costs, except to the extent that the buyer is able to take the car to a mechanic and have it inspected prior to purchase. In any event, high maintenance costs do not seem to have any significant effect on the resale price of a vehicle.

A final variable, the vehicle’s internal rate of return (IRR) over its lifetime, is a significant predictor of resale values, at least for the compact and RV. One interpretation is that a vehicle with a high IRR represents a “good car” that is attractive and was frequently rented by customers (thus resulting in the high IRR). However the interpretation of this coefficient is problematic due to potential endogeneity of IRR. As noted in table 0, a car with a higher resale value will have a higher IRR, all other things equal. Thus, unobservable characteristics of a car that lead it to have a higher resale value could also lead it to have a higher IRR, and thus the positive coefficient on IRR could be partly spurious, due to a positive correlation between IRR and unobservable factors affecting a vehicle’s resale value.

Overall, the main conclusion we draw from table 1 is that beyond age and odometer value (and implicitly the car’s characteristics, as represented by its make and model), there are few other significant explanatory variables for the resale value of a car. Our regressions can explain only between 40 to 50% of the variation in the resale values of the cars the company sells: there is a lot of “residual variance” that leads one car to sell for much more than another car that from our standpoint is “observationally equivalent” to it, at least in terms of the variables we can observe. To the extent that this residual variation in used sales prices really is “random” and cannot be accounted for by unobserved idiosyncratic characteristics that the buyers can see but which we do not have in our dataset, there may be other room for this firm to increase its profits by adopting a “reservation price strategy” when it sells its cars. That is, instead of putting a car up for sale (in some cases in an auction), and selling the car no matter what, even if the best offer for it is very low, if the firm adopts a reservation price strategy it would refuse to sell the car on a particular occasion if the best offer price was not sufficiently high. Instead, it would try to resell the car at a subsequent date, perhaps in a different venue. In the analysis below, we do not assume that there is anything suboptimal about the selling strategy that this firm uses. We assume that the large variation in resale prices is justified based on unobservable characteristics of the cars they are selling that we cannot observe (e.g. scratches or dents, stains or rips in the upholstery, etc.) that account for the large variation in resale prices among "observationally equivalent" vehicles that we observe. Later in section 6 we will consider how our conclusions would be altered if we adopt the alternative hypothesis that a good deal of the variation in resale prices is not justified, and represents some sort of inefficiency or variability in bids for cars up for sale, that could signal inability of the firm to achieve the "fair market value" for some of the cars it sells.

Now we consider the econometric estimation of patterns of usage of vehicles during rental spells. As we noted above, the firm does not impose a maximum number of kilometers that can be driven during a rental contract or any per kilometer usage rate, and so rental customers are entirely free in their choice of how much to drive their rental cars during a contract. The level of utilization by rental customers is obviously a key part of the "law of motion" because it determines how quickly a car will "age" in terms of the car's odometer value, which we demonstrated above is a key predictor of the resale value of the car. However there is a difficulty, noted above, is that the firm frequently does not accurately record the in and out odometer values for its vehicles, making it impossible 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 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 travelled each day by a rental customer are *IID* 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.[6] The probability density function corresponding

[6]Actually, 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 $\alpha = d$.

to F is given by

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

where $\Gamma(d)$ is the *Gamma function.* Thus, $o'-o$ is the actual number of kilometers travelled 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 travelled 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 travelled 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 travelled per day for short term and long term contracts separately. However since the odometer values in and out of rental contracts do not appear to be accurately recorded in the company's data, we cannot use this approach. Indeed, if we had such data, it might even be possible for us to estimate the conditional distributions $F(o'|o,d,r)$ non-parameterically or semi-parametrically, and thus, not have to rely on the parametric assumption that kilometers travelled per each day of a rental contract are *IID* exponential variables. However we can estimate the two λ_r parameters necessary for us to determine the distributions $F(o'|o,d,r)$ since we do have accurate records 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 at time of sale is given by

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

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

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

Since we do accurately observe the number of days a vehicle is rented, 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.

Using these regression estimates and the information on (N_t^l, N_t^s) for each car, now indexed by the day in its service life, t, we can 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 our predicted odometer values are reasonably accurate. Using these predicted odometer, our next step is to analysis 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

Table 3 presents the results of a binary logit model of the company's selling decision for the three car types, compact, luxury, and RV. We tabulated data for each day a car was in a lot spell and treated the company as having the opportunity to either keep the car or sell it on that day. Thus, if we let s_t denote a binary variable for the selling decision with $s_t = 1$ if the company sells the car and $s_t = 0$ if the company keeps the car, we estimated the parameters θ as coefficients of variables x_t that represent different factors that might affect the company's decision to sell the car, using the standard logistic functional form

$$Pr\{s_t = 1 | x_t\} = \frac{\exp\{x_t\theta\}}{1 + \exp\{x_t\theta\}}. \quad (8)$$

Among the variables in the vector x_t are the vehicle's age and predicted odometer value (based on the regression estimate $\hat{o}_t$ using the observed values of N_t^l and N_t^s from the rental contract data, as discussed above), the duration in the lot, the average daily maintenance costs for the vehicle at day t (the ratio of the total maintenance costs incurred up to date t divided by t), and the vehicle's utilization rate (the total number of days the vehicle was rented up to t divided by t).

The empirical results here are somewhat mixed. As we previously discussed, due to the collinearity between age and odometer value, it is difficult to identify the separate effects of age versus odometer value on the firm's decision to sell a vehicle. For the compact and RV, the age variable is the statistically significant predictor and odometer value is statistically insignificant. However for the luxury car, the opposite result holds: odometer is the statistically significant predictor and age is insignificant. The overall goodness of fit of the models does not change significantly if we use only age or only odometer values to predict the company's selling decision.

	Compact All Locations		Luxury All Locations		RV All Locations	
Variable	Estimate	(t-stat)	Estimate	(t-stat)	Estimate	(t-stat)
Constant	-13.061	(-22.34)	-12.271	(-5.41)	-14.669	(-8.39)
Age (days)	0.0077	(7.34)	0.0011	(0.43)	0.0125	(5.82)
Odometer (km)	0.0050	(0.44)	0.0987	(3.37)	-0.038	(-1.91)
Duration, Age < 500	0.0206	(12.31)	-11.994	(-19.23)	-6.069	(-49.78)
Duration, Age $\in [500, 1000)$	0.0867	(24.37)	0.0471	(8.85)	0.0399	(7.72)
Duration, Age > 1000	0.1362	(13.23)	0.1736	(2.85)	0.1744	(6.49)
Maintenance Cost	0.00003	(2.14)	0.2030	(0.85)	-0.0188	(-0.23)
Utilization Rate	0.4049	(0.61)	-1.616	(-0.57)	1.989	(0.84)
N, $\log(L)/N$	36262	-0.017	6445	-0.022	7192	-0.017

Table 3: Logit Estimation Results for Decision to Sell Car

Besides age and odometer, the only variable whose coefficient estimates are statistically significant and has signs that are (generally) consistent with our *a priori* expectations is the duration variable. For young cars (less than 500 days old), duration has a negative coefficient for the luxury and RV cars, and a positive coefficient for the compact. For cars that are less than 500 days old, there is very little chance that they will be sold, and so the duration variable is mainly acting as a "dummy variable" and the negative values are the model's way of telling us that the chance of being sold is close to zero. However when the car is between 500 and 1000 days old, all three duration coefficients are positive and are uniformly higher than in the case for cars less than 500 days old. Finally for older cars, i.e. those whose age is greater than 1000 days, the duration coefficients are the highest. This tells us that the cars that are at the greatest risk of being sold are the older cars that have been on the lot unrented for many days. Besides this strong duration effect, neither maintenance costs nor utilization rates (or any other variables we considered in estimations that we have not reported due to space constraints) appear to have any strong and consistent effect on the company's decision to sell their vehicles.

Figure 4 summarizes the main factors governing the replacement decisions by this firm for each of the three vehicles. The left hand column of figure 4 plots the cumulative distribution function for replacements as a function of the odometer value, and the right hand column plots the cumulative distribution function in terms of the vehicle age. The right hand column confirms the company's claim that its targeted 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 column 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 more on the "3 year rule" than on a rule based on number of kilometers driven. However an alternative hypothesis is that the company makes replacement decisions on an odometer threshold, but that it has different "optimal" thresholds for different types of vehicles. We will see that while it is hard to empirically distinguish which of these two hypotheses is the "correct" it is not really necessary to make such a distinction for our purposes. 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.

We do note that the firm is clearly not following an *exact* age or odometer *threshold replacement rule.* That is, the firm does not replace all of its cars at the moment they exceed three years of age, or the instant their odometer values exceed some specified cutoff value. Instead, we see a fairly wide range of ages and odometer values over which cars are sold. The oldest ages at which cars are sold occurs at 4, 4.9 and 3.8 years, repectively for the three types of vehicles, where the earliest ages at which vehicles are sold is at 1.5, 1.9 and 1.9 years respectively. Thus, replacements occur over 2 to 3 year interval starting around 2 years of age, by 2.8 years 50% of the cars have been replaced, and after 5 years of service all of the cars have been sold off. In terms of odometer values, the range is even larger: one compact car was sold with only 10,000 kilometers on its odometer. The highest odometer at which we observed any of these cars being sold was a luxury vehicle that was sold at 187,000 kilometers. The relatively small numbers of relatively "old" cars that the firm sells lead to distributions of ages and kilometers at sale that are skewed to the right. Similar to predicting resale prices, we conclude there are a number of unobservable factors that lead to relatively wide variation in odometer values and ages at which different cars are sold. Besides age and odometer value (and secondarily duration in the lot, when cars are sufficiently old), we have not found any other variables in our database that are capable of predicting why certain cars are sold when they are relatively young and others are sold when they are relatively old.

The remaining objects to be estimated to implement our econometric model 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 choose 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, $d+1$. Thus with probability $1-h(d,r)$ it

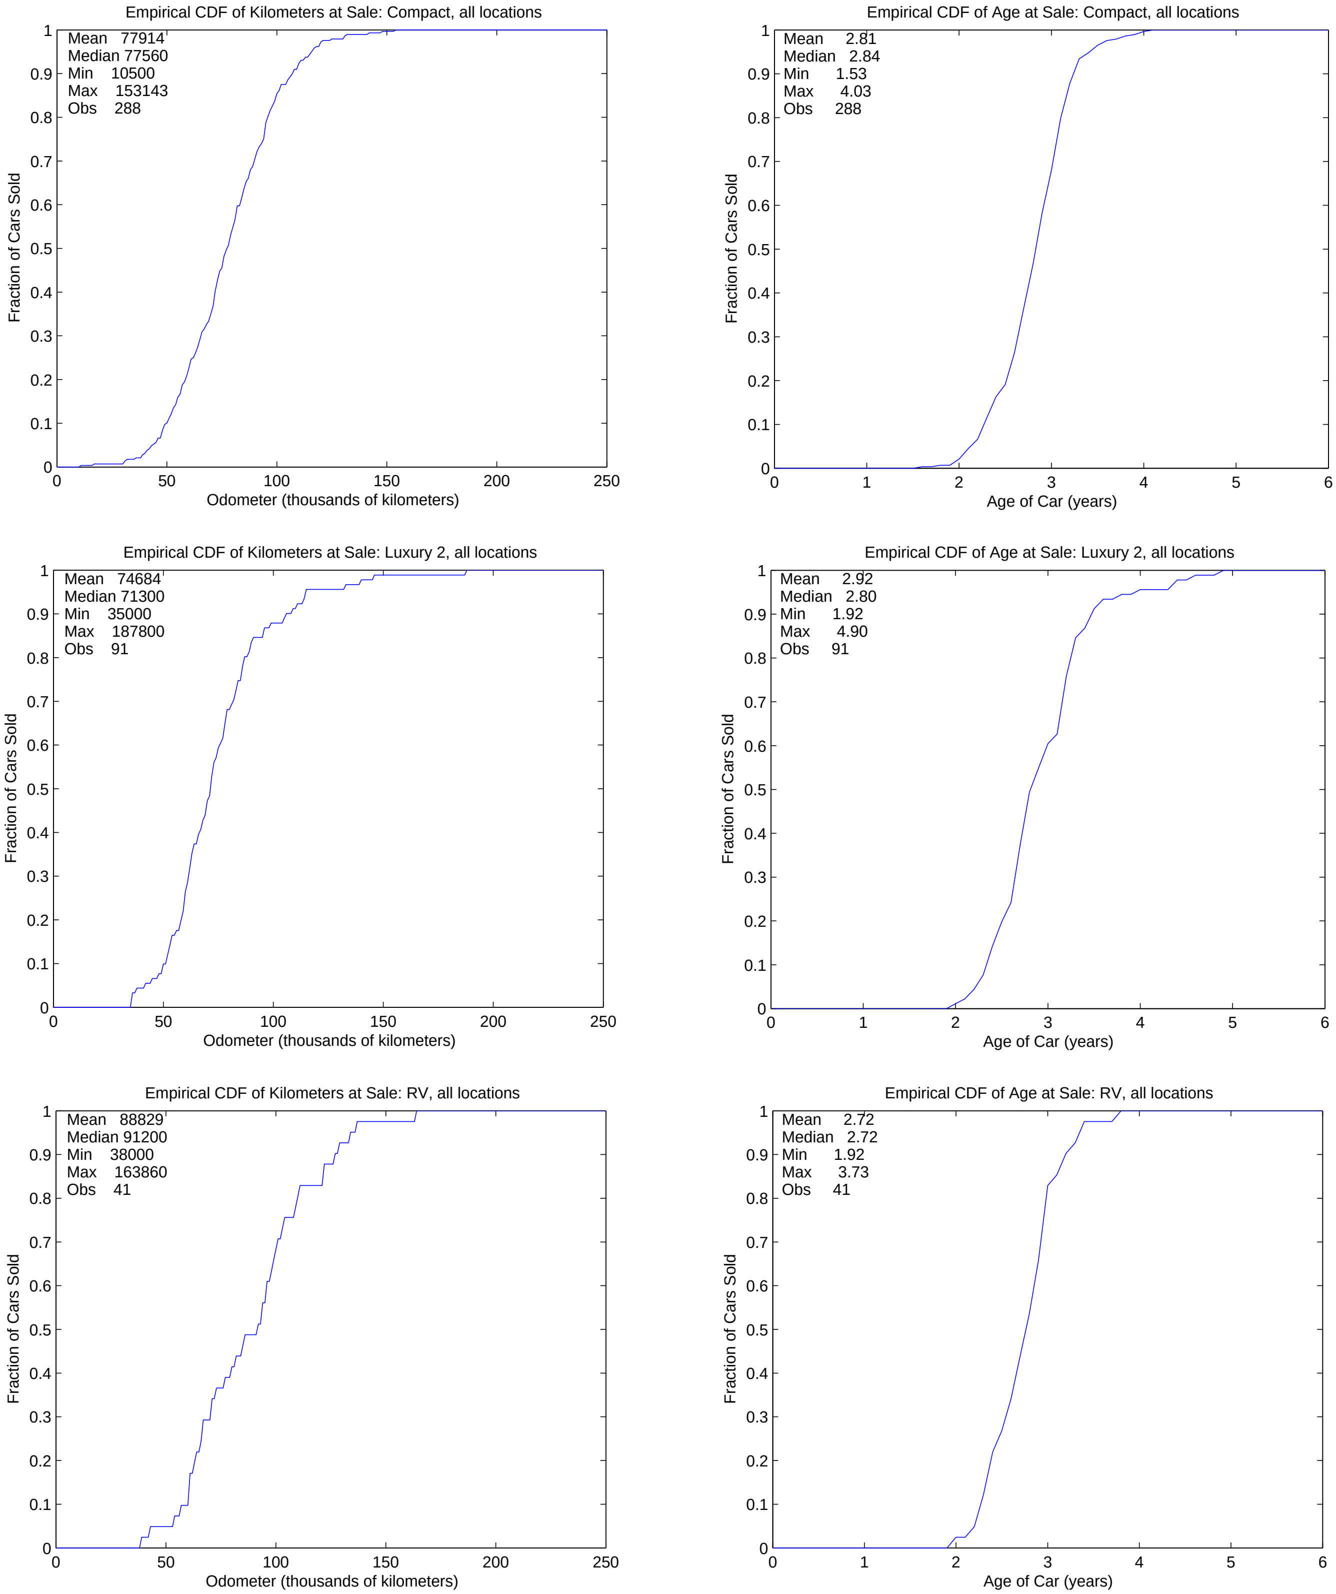

Figure 4 Cumulative Fractions of Cars Replaced: Compact, Luxury, and RV – All Locations

will continue to remain in state r. The following well known recursion formula allows us to compute the duration distribtion $f(d|r)$ implied by the hazard function $h(d,r)$

$$f(d|r) = \begin{cases} f(1|r) = h(0,r) \\ f(d|r) = \prod_{j=0}^{d-2} [1 - h(j,r)]\, h(d-1,r) \quad d \geq 2 \end{cases} \tag{9}$$

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. Very few long term rentals last fewer than 15 days, perhaps in part due to the 20% penalty the company imposes on early return of vehicles in a long term contract. Figure 5 presents the duration distributions implied by our estimated hazard rates.

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, our 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 tail hazard rate tail 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 or 4 lot spell. Figure 5 shows the rental duration distributions (left panel) and the lot spell hazard functions (right panel). The duration distributions and lot hazard functions for the other two car types are similar, and are omitted due to space constraints. Note that none of our estimated hazard rates for the various spell durations depend on the vehicle's odometer. We have omitted the odometer variable since there is no evidence for any "aging effect" in the durations for any of the 4 possible spell types. To convince the reader, figure 6 presents scatterplots of the durations for all spells for all of the compact cars in the data set.

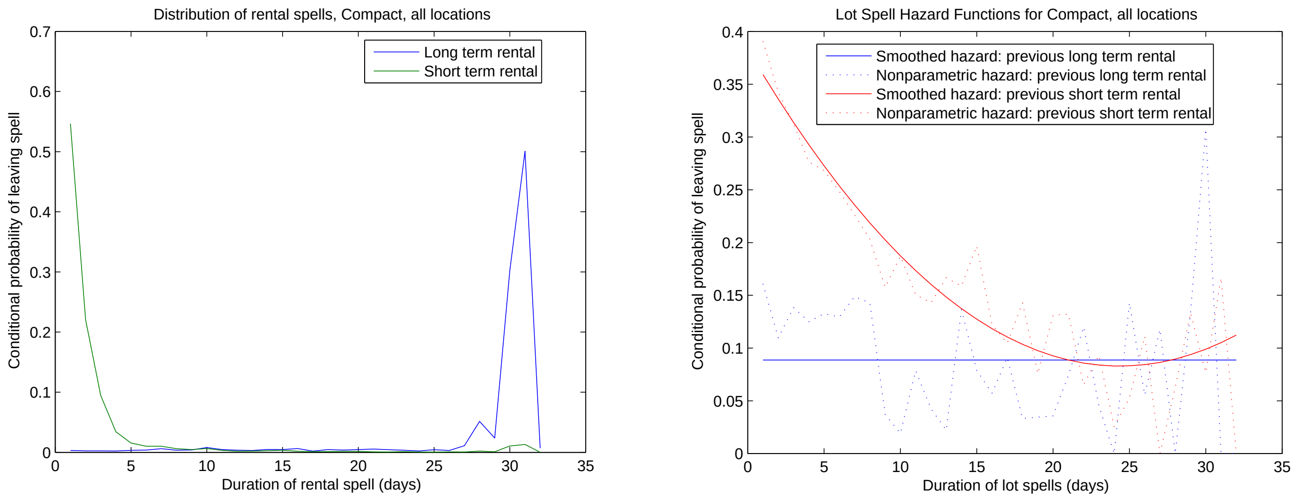

Figure 5 Rental Spell Duration Distributions and Lot Spell Hazard Functions: Compact – All Locations

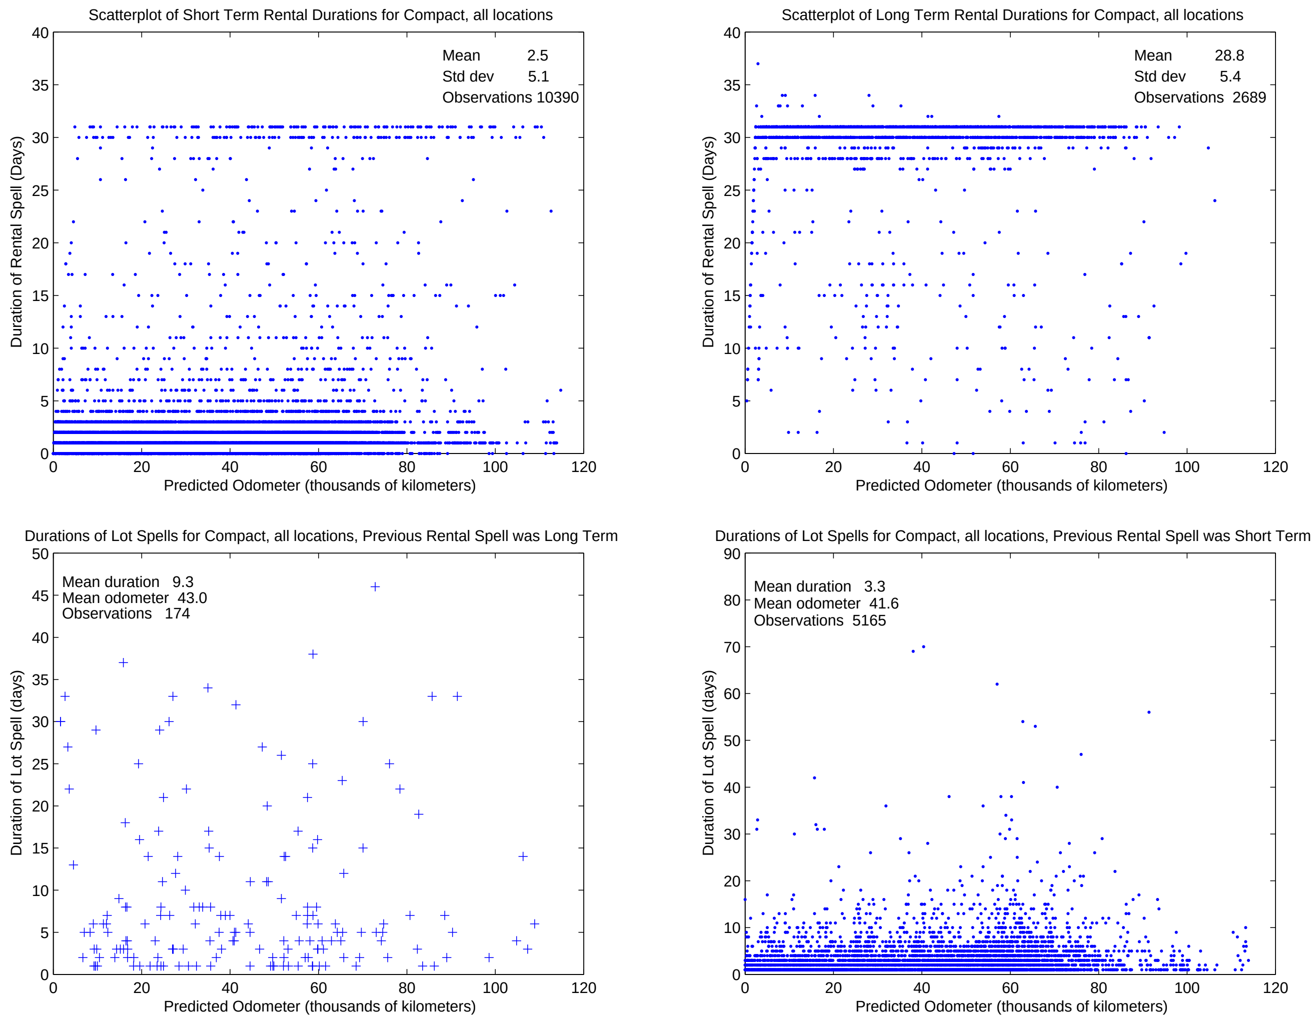

Figure 6 Scatterplots of Lot and Rental Spell Durations: Compact – All Locations

The x-axis in these scatterplots is the predicted odometer value. The results would be essentially unchanged if we had used the age of the vehicle in days (for which no regression predictions are necessary since in and out dates for rental spells are accurately recorded in the database). We see that there is no trend in spell durations as vehicles age, thus justifying our decision not to include the vehicle's odometer value as a covariate in the estimated hazard rates, $h(d,r)$.

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 probability 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.* If $r > 2$, i.e. the car is in a lot spell, we rule out "self transitions", i.e. $\pi(r|r,d,o) = 0$ for $r > 2$. This is 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, perhaps the previous customer renewing or extending 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), and 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 evident from the right hand panel of figure 5: 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 thus mean durations are higher. In plain language, if a car had previously been in a long term rental and the contract did not immediately roll over, one can 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 rental spells (i.e. 1) long term contract, 2) short term contract, or 3) lot spell), we used a trinomial logit model to estimate these probabilities. This probability is given by

$$\pi(r'|r,d,o) = \frac{\exp\{v(r,d,o)\theta_{r'}\}}{\sum_{\rho\in\{1,2,l(r)\}}\exp\{v(r,d,o)\theta_\rho\}}, \tag{10}$$

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. Table 4 presents the results of the trinomial logit estimation for the three car types.

	Compact All Locations		Luxury All Locations		RV All Locations	
Variable	Estimate	(t-stat)	Estimate	(t-stat)	Estimate	(t-stat)
θ_2: coefficients for transition to short term rental						
Constant	4.602	(27.45)	3.012	(11.01)	4.137	(10.68)
Odometer, o (000 km)	0.0113	(4.24)	0.002	(0.69)	0.001	(0.27)
Duration, d	−0.068	(−6.29)	−0.039	(−2.92)	−0.087	(−3.86)
$I\{d >= 29\}$	−0.421	(−1.52)	0.006	(0.02)	0.079	(1.27)
$I\{r = 1\}$	−6.653	(−31.20)	−6.289	(−20.72)	−6.326	(−13.77)
$\theta_{l(r)}$: coefficients for transition to lot spell of type $l(r)$, where $l(1) = 2$ and $l(2) = 4$						
Constant	3.879	(23.17)	3.696	(14.29)	4.359	(11.38)
Odometer, o (000 km)	0.0204	(7.67)	0.007	(2.52)	0.010	(2.25)
Duration, d	−0.077	(−7.81)	−0.082	(−7.95)	−0.120	(−6.31)
$I\{d >= 29\}$	−1.504	(−6.79)	−1.076	(−4.91)	−0.775	(−1.76)
$I\{r = 1\}$	−4.491	(−25.32)	−3.442	(−16.43)	−3.775	(−11.75)
N, $\log(L)/N$	16246	−0.606	3617	−0.484	2142	−0.583

Table 3: Trinomial Logit Estimates of Transitions out of Rental Spells

For transitions out of lot spells, since we have ruled out the possibility of “self-transitions” there are only two possible destinations: long term rental spells and short term rental spells. Table 4 presents binary logit estimation results for transitions from lot spells of type 3 and 4 (i.e. where the previous rental spell was a long term contract and short term contract, respectively), i.e. for transition probabilities specified as

$$\pi(r' = 1|r,d,o) = \frac{\exp\{v(o,d)\theta_r\}}{1+\exp\{v(o,d)\theta_r\}}, \quad r \in \{3,4\} \tag{11}$$

	Compact All Locations		Luxury All Locations		RV All Locations	
Variable	Estimate	(t-stat)	Estimate	(t-stat)	Estimate	(t-stat)
θ_3 coefficients for transitions from type 3 lot spells						
Constant	2.262	(3.94)	1.590	(3.53)	2.335	(2.09)
Odometer, o (000 km)	0.010	(0.83)	−0.002	(−0.23)	−0.009	(−0.59)
Duration, d	−0.053	(−2.39)	−0.004	(−0.45)	−0.038	(−1.05)
N, $\log(L)/N$	173	−0.326	181	−0.490	43	−0.511
θ_4: coefficients for transitions from type 4 lot spells						
Constant	3.634	(16.86)	1.935	(8.97)	4.536	(9.67)
Odometer, o (000 km)	0.0211	(3.74)	0.013	(2.51)	−0.009	(−1.12)
Duration, d	−0.064	(−4.46)	−0.003	(−1.10)	−0.026	(−2.28)
N, $\log(L)/N$	5162	−0.077	961	−0.683	922	−0.090

Table 4: Binomial Logit Estimates of Transitions out of Lot Spells

There are two key points to take away from tables 3 and 4: 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 and the probability of transitions into short term rental contracts increases.

Figure 7 illustrates these age effects, i.e. it plots $\pi(r'|r,d,o)$ as a function of o for fixed (r,d). The two top panels plot the three possible probabilities in a transition from the 31st day of a rental spell (i.e. $d = 31$) for $r = 1$ (previous spell was a long term rental, top left panel of figure 7) and for $r = 2$ (previous spell was a short term rental, top right panel of figure 7). We see that for new cars, there is a very high probability of a transition into a long term rental contract, but as the car ages, this probability falls. The bottom two panels of figure 7 show the estimated transition probabilites into a long term rental contract from a lot spell of type 3 (bottom left panel of figure 7) and from a lot spell of type 4 (bottom right panel of figure 7).

The bottom panels of figure 7 plot three probability curves, for durations of $d = 1$, $d = 10$ and $d = 30$ in the lot spell. We see a strong duration effect (i.e. probability of transition into a long term rental spell increases with duration in the lot spell), and strong aging effects (i.e. the probability of transition into a long term rental spell, for any duration, decreases with the odometer value of the car). Overall, the age effects shown in figure 7 constitute the "rental contract composition age effect" that we discussed in the introduction. Simply put, as the car ages, whenever it is rented, the more likely it is that the rental will be a short term contract. Further, the likelihood of being in a lot spell also increases with the age of the vehicle.

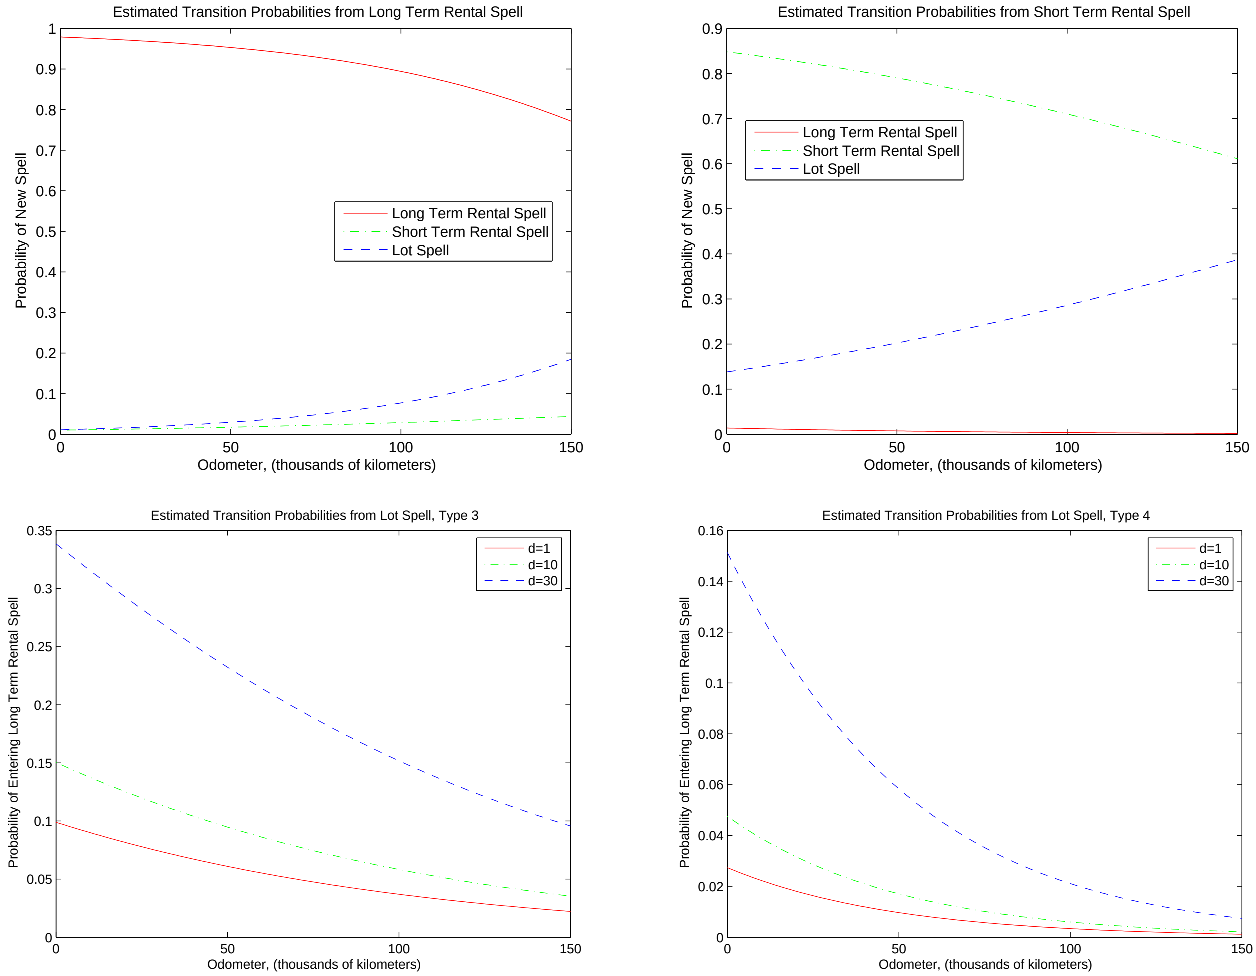

Figure 7 Estimated Transition Probabilities as a Function of Odometer Value: Compact – All Locations

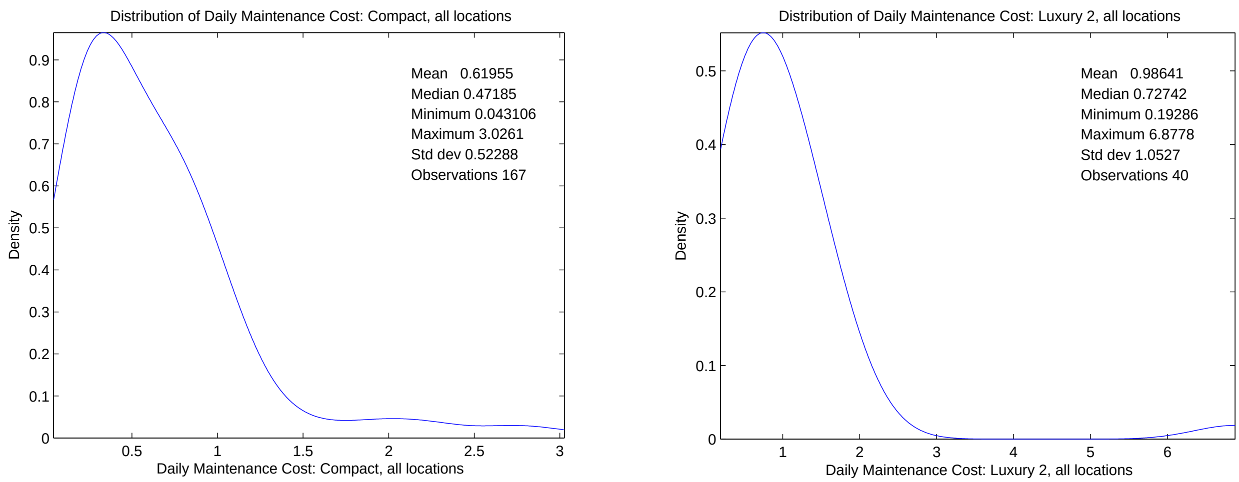

Figure 8 Distribution of Daily Maintenance Costs: Compact and Luxury – All Locations

The final objects necessary to complete our econometric model are the daily rental rates and maintenance costs. The estimation of daily rental rates is trivial: in fact since the rates are fixed, we can use the company's published tariff rate. However there is some minor variation in daily rental rates due to variations in optional equipment and features on cars (e.g. some cars have larger engines that 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 but the number of days in these spells.

Maintenance costs are incurred on a 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 occuring 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 events, and then conditional on an event occurring model the distribution of maintenance costs occurred on each maintenance occasion, 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.

Figure 8 plots the distribution of daily maintenance costs for the compact and luxury car types, respectively (the distribution of maintenance costs for RVs is similar and not shown due to space constraints). We see that there is evidence of some car-specific heterogeneity in total maintenance costs, with evidence that a minority of cars are "outliers" or "lemons" in the sense that their daily maintenance costs are 5 to 6 times higher than the average. However it appears that high maintenance costs are a relatively "second order" consideration for this company in terms of vehicle replacement decisions. Recall that daily maintenance costs did not emerge as a significant predictor of either used car prices, or of the company's decision to replace a car. As a result, we did not feel that the payoff from a more detailed model of the timing of maintenance would be very high. Although charging an average daily maintenance costs is a less accurate way of modeling the actual incidence of cash flows, since maintenance costs are small relative to rental revenues, the discounted value of profits and other measures such as the internal rate of return on the vehicle are not sensitive to the timing assumptions. Indeed, in the next section we will now show, via the stochastic simulations, that our simplified treatment of maintenance costs does not compromise our ability to provide a good overall model of the firm's operations.

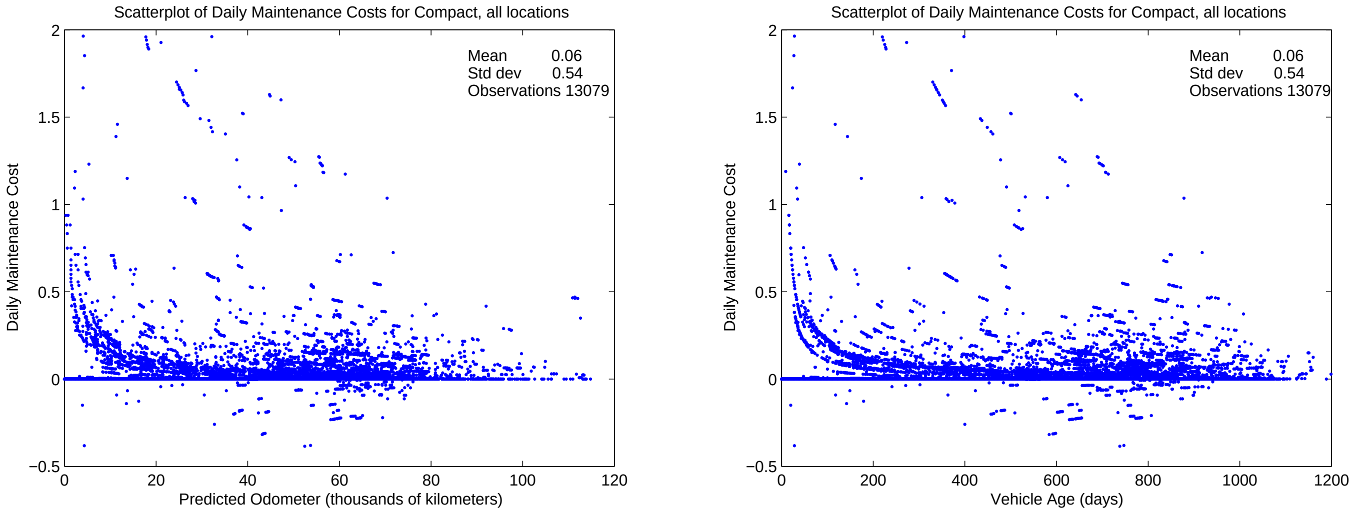

Figure 9 Scatterplot of Average Daily Maintenance Costs in Previous 30 Days: Compact – All Locations

The final issue we address is whether there is evidence of aging effects in vehicle maintenance costs. Figure 9 presents a scatter plot of the average daily maintenance costs incurred by the company on its compact vehicles as a function of the predicted odometer value. The figure shows no evidence that these daily costs increase with the odometer value of the vehicle. The results are essentially unchanged if we plot daily maintenance costs as a function of the vehicle age (right hand panel of figure 9). Thus, we conclude that the only aging effects that we can detect in our 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 tree vehicle types analyzed in section 3.

Our stochastic simulation program starts with the purchase of a new vehicle. Random numbers are drawn from a uniform random number generator to determine various outcomes. In general the "probability integral transform method" is used to construct random draws from various conditional distributions.

For example, assuming new cars start out in a lot spell, let $F(d)$ denote the duration distribution for the initial lot spell. A random duration $\tilde{d}$ in that spell is constructed by taking the inverse cumulative distribution evaluated at a random uniform draw $\tilde{u} \sim U[0,1]$. That is, we draw $\tilde{u}$ and compute $\tilde{d} = F^{-1}(\tilde{u})$, where F^{-1} is the inverse of the cumulative distribution function, defined by

$$F^{-1}(u) = \sup_{d}\{F(d) \leq u\}. \tag{12}$$

If F is strictly increasing (which is true for continuous random variables) then F^{-1} is the usual inverse function satifying $F(F^{-1}(u)) = u$ and $F^{-1}(F(d)) = d$. However the general definition is also valid in cases where $\tilde{d}$ is a discrete random variable, as in the case where $\tilde{d}$ is the duration of a spell measured in days. We rely on the well known fact that if $\tilde{u} \sim U[0,1]$, then $\tilde{d} = F^{-1}(\tilde{u}) \sim F$, i.e. $\tilde{d}$ will have the correct distribution function F.

At the end of the initial lot spell, the simulated car will make a transition into either a long term rental contract or a short term rental contract with probability $\pi(1|\tilde{d}, r_0)$ and $1 - \pi(1|\tilde{d}, r_0)$, respectively, where π is the rental state transition probability that we introduced (and estimated) in section 3, and r_0 denotes the initial lot spell. We determine the rental state by drawing a $\tilde{u} \sim U[0,1]$ random variable and if it is less than $\pi(1|\tilde{d}, r_0)$ the car makes a transition to a long term rental spell, otherwise it enters a short term rental spell. The duration in a rental spell is simulated in the same way. At the end of the rental spell, another $\tilde{u} \sim U[0,1]$ random variable is drawn in order to generate the number of kilometers driven during the rental spell. We use the Erlang distribution $F(o'|0,d,r)$ described in section 3, and generate the "in odometer" reading (i.e. the odometer on the car when it returns from the rental spell) via the probability integral transform, $\tilde{o}' = F^{-1}(\tilde{u}|0, \tilde{d}, \tilde{r})$, where $o = 0$ is the initial odometer reading on a brand new car, and $\tilde{d}$ is the simulated duration of the rental spell, and $\tilde{r}$ is the simulated rental spell type (long or short term).

This process continues until the point where a simulated vehicle replacement occurs, at which point the simulation of this vehicle terminates and a random resale price is determined from a lognormal distribution for resale values implied by the lognormal regression model for vehicle depreciation rates estimated in section 3. Specifically, we draw a random resale price $\tilde{P}_t(o_t, \tau)$ for car type τ after a sale on day t using a randomly drawn standard normal random variable $\tilde{x} \simeq N(0,1)$ and the equation

$$\tilde{P}_t(o_t, \tau) = P_0(\tau) \exp\left\{\hat{\alpha}_1(\tau) + \hat{\alpha}_2(\tau) + \hat{\sigma}(\tau)\tilde{x}\right\}, \tag{13}$$

where $(\hat{\alpha}_1(\tau), \hat{\alpha}_2(\tau), \hat{\sigma}(\tau))$ are the estimated lognormal regression coefficients and standard error for car of type τ from the logarithmic price depreciation regression in equation (3) in section 3.

We simulate the company's decision to sell a vehicle based on the vehicle's odometer value, using the cumulative replacement probability distributions in the left column of figure 5. Thus, at the start of each simulation, we draw a random "replacement threshold" $\overline{o}$ from the cumulative distribution function $G(o|\tau)$, where G is the cumulative probability that a car of type τ with odometer value o or lower will be sold (the distributions in the left hand column of figure 5). Then as the simulation proceeds, the car will be sold at the first eligible instance where the simulated rental vehicle's odometer exceeds the random selling threshold $\overline{o}$. We use the term "first eligible instance" since we respect the fact that the company will not sell a car if it is in the middle of a rental contract. We assume that the only feasible time for the company to sell a car is the day it returns from a rental spell, or any day that the car is in a lot spell. In the former case, if the firm decides to sell vehicle that a customer wants to rent again, it would provide the customer with a newer replacement.

Figure 10 shows individual simulated rental histories for each of the three car types that we analyze, i.e. compact, luxury, and the RV. Comparing this to figure 1, we see that the simulated histories appear quite similar to actual rental histories for the three sample vehicles presented in figure 1. Of course, there is a vast number of possible rental histories (indeed infinitely many of them) and the fact that we can find three simulated histories that resemble the three actual rental histories we chose to display in figure 1 should not necessarily convince a skeptical reader that our model is an adequate one. We need to show that our estimated semi-Markov model of the company's rental operations is capable of capturing *the entire distribution of possible outcomes* for a wide variety of different measures of the operations, revenues, profits, and returns earned by this firm. If we can show that our model provides a sufficiently close approximation to the entire range of actual outcomes this company experiences, the reader can have more confidence that our econometric model is a good one, and that it might be of use for simulating how profits, rates or return and other variables of interest would change if *alternative operating strategies* were adopted.

Figure 11 presents comparisons between simulated and actual distributions of the odometer and vehicle age (in days) at which vehicles are replaced. Due to space constraints we present the results for the luxury car type only, although the results for the compact and RVs are similar. The left hand panel of figure 11 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 since, by construction, we have drawn the odometer values at which the company replaces its vehicles from the actual (empirical) distribution. Thus, the differences in the two distributions in the

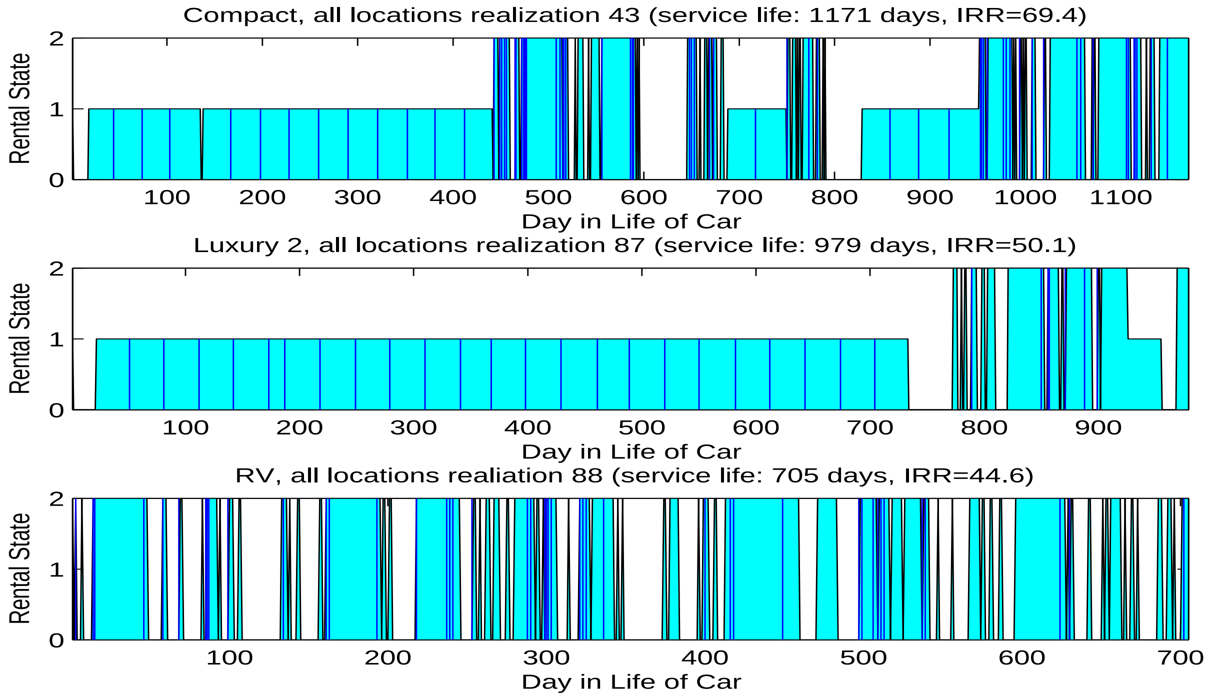

Figure 10 Simulated Rental Histories for the Three Car Types, Compact, Luxury and RV

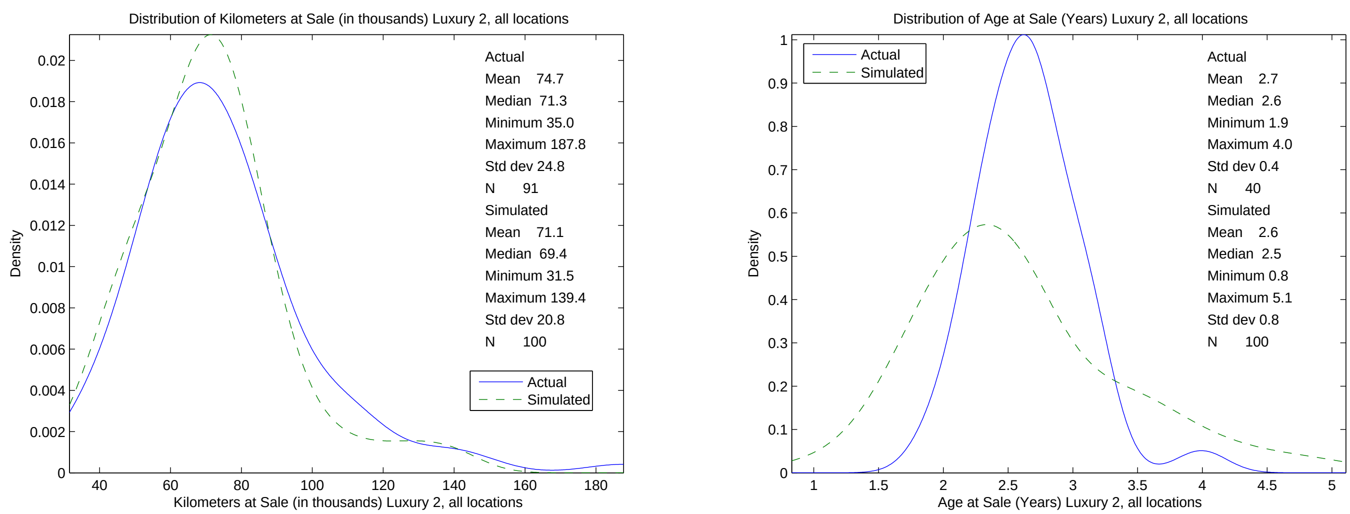

Figure 11 Simulated versus Actual Odometer and Age at Sale: Luxury – All Locations

left hand panel of figure 11 are entirely due to *sampling error* in our random sample of 100 simulated cars.

The right hand panel of figure 11 compares the actual distribution of replacement ages to the one implied by our 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 (see middle right panel of figure 4), 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 our 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 random replacement threshold $\overline{o}$, there is a an implied replacement age as well.

The simulated and actual distributions of ages at replacement are further apart than for the distributions of odometer values at replacement, although we do note that the mean simulated age at replacement, 2.6 years, is very close to the actual value, 2.7 years. The simulated distribution of replacement ages does have 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 more closely based on age than odometer value, and thus more tightly concentrated around the three year replacement target that we discussed in previous sections. 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, with cars that have high simulated utilization rates being younger than average at time of sale, whereas those with low simulated utilizations rates being older than average at time of sale.

While we could adopt a more complex replacement rule that is based both on age and odometer and other variables such as duration in the lot spell before the vehicle was replaced, we feel that the simpler odometer-based replacement rule provides a sufficiently good approximation to the company's behavior and outcomes — as we will see in figures 12 and 13 which compare 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 as we shall see in the next section, this enables us to cast the problem into a stationary Markovian decision process formulation and thus to estimate the expected present discounted value of the company's profits over an infinite horizon. We will 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.

Figure 12 shows that our econometric model provides a good approximation to the distribution of different rental outcomes over a vehicle's service life. The left hand panels compare the simulated versus actual distributions of the number of long term and short term rental spells, and the number of lot spells. The right hand panels compare simulated versus actual distributions of total days spent in each of these spells. We see that our model not only does a good job of matching the mean values of the number of spells and durations of each spell type, but it also does an extremely good job of capturing the overall distributions as well. Matching both the number of duration of the various types of spells turns out to be the key to accurate predictions of revenues, profits and returns.

Figure 13 plots comparisons of simulated versus actual distributions of the relevant financial variables. The top left panel of figure 13 shows that our econometric model results in a distribution of proceeds from sales of cars that is quite close to the actual distribution. This is evidence that our lognormal regression model of vehicle price depreciation is a good one. The top right hand panel of figure 13 compares the actual and simulated distributions of total maintenance costs. Although the mean maintance costs are values are close to each other, the simulated distribution of maintenance costs has much less variance. This is to be expected, given our "shortcut" procedure for simulating maintenance costs that we described in the previous section. That is, instead of trying to estimate the durations between successive maintenance episodes, we simply "smoothed out" the distribution of maintenance costs to an equivalent per day maintenance cost, and thus, our estimate of total maintenance costs is simply equal to 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, as we will see shortly, 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 $1/12$th of the average resale value for the luxury car type, and about $1/23$rd of the cost of a new car.

The middle two panels of figure 13 plot the distribution of total revenues earned from long and short term rental contracts. Once again, our econometric model does an extremely good job not only in matching the mean revenues, but also in capturing the overall distribution of revenues. We see that for the luxury car type, 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,

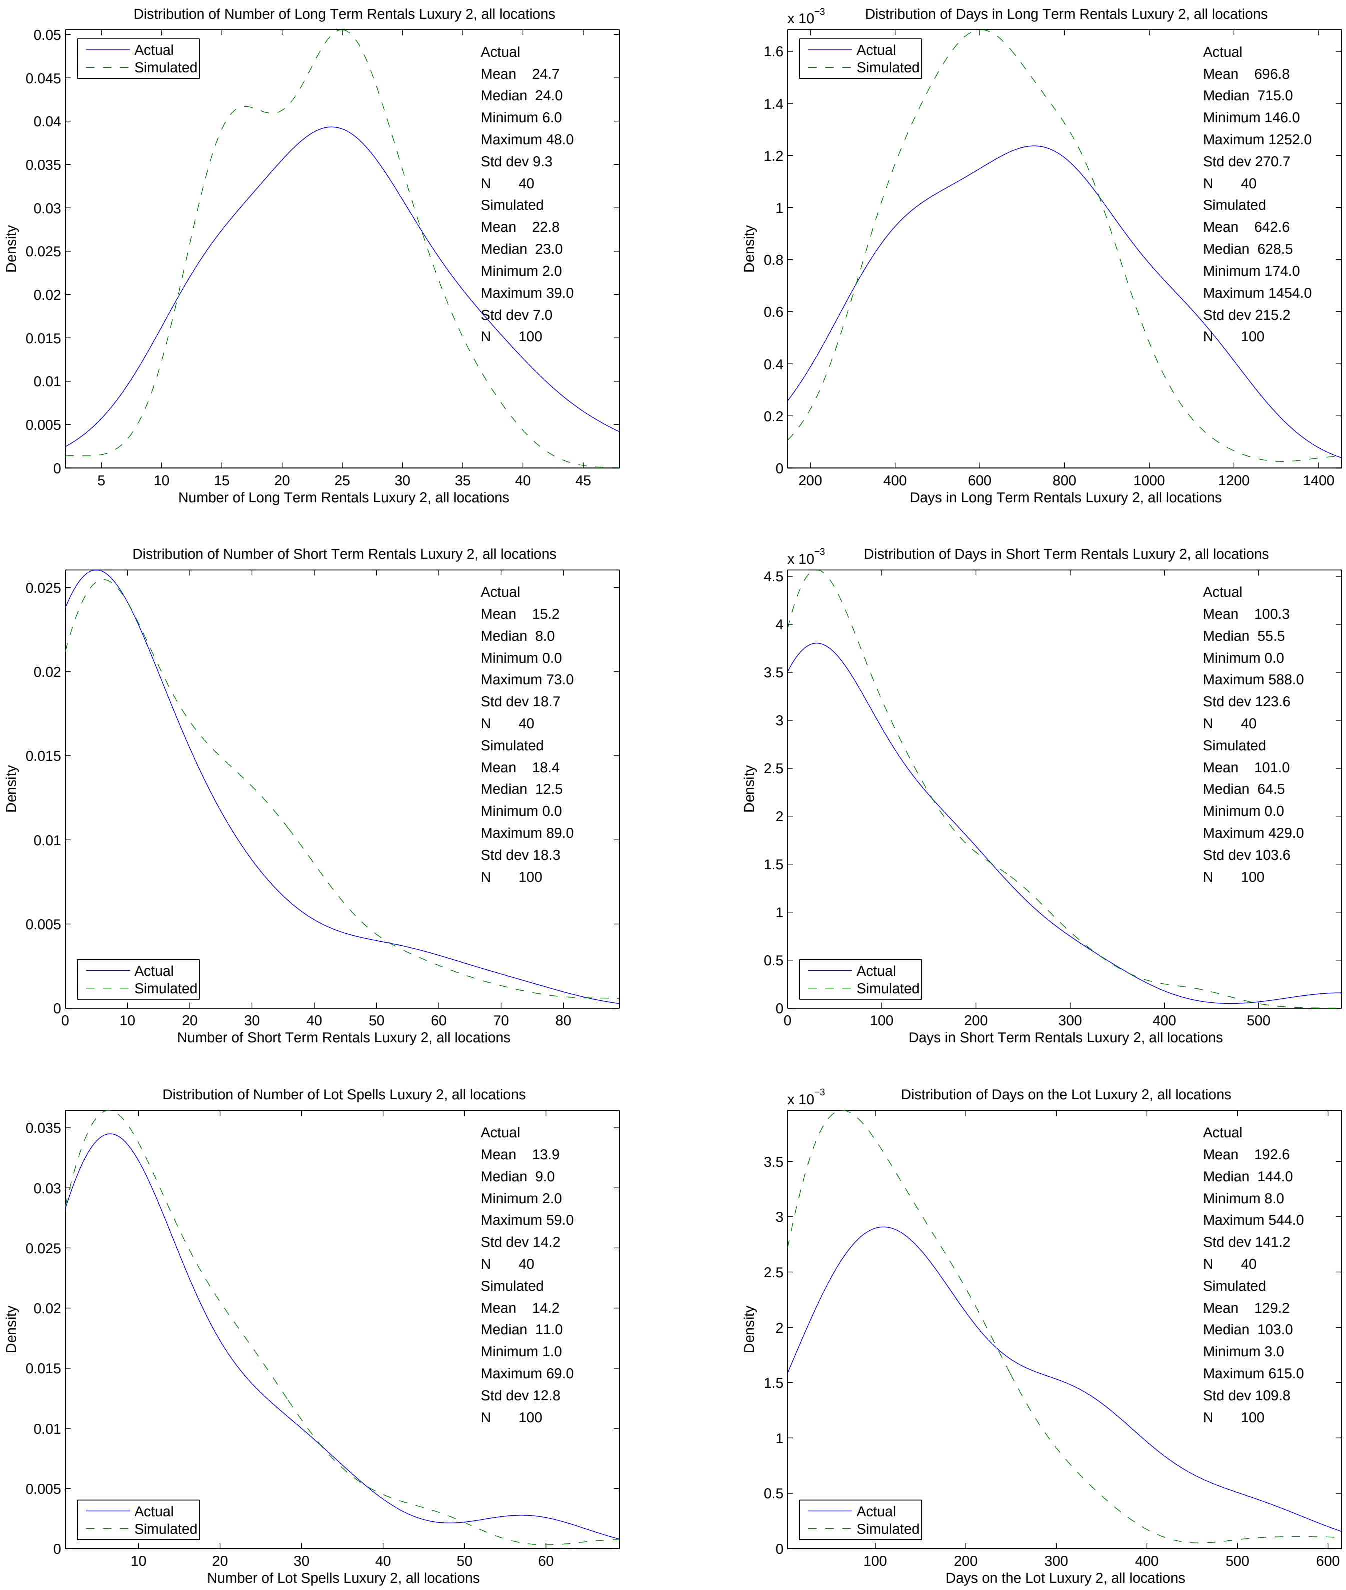

Figure 12 Simulated versus Actual Number of, and Days in Spells: Luxury – All Locations

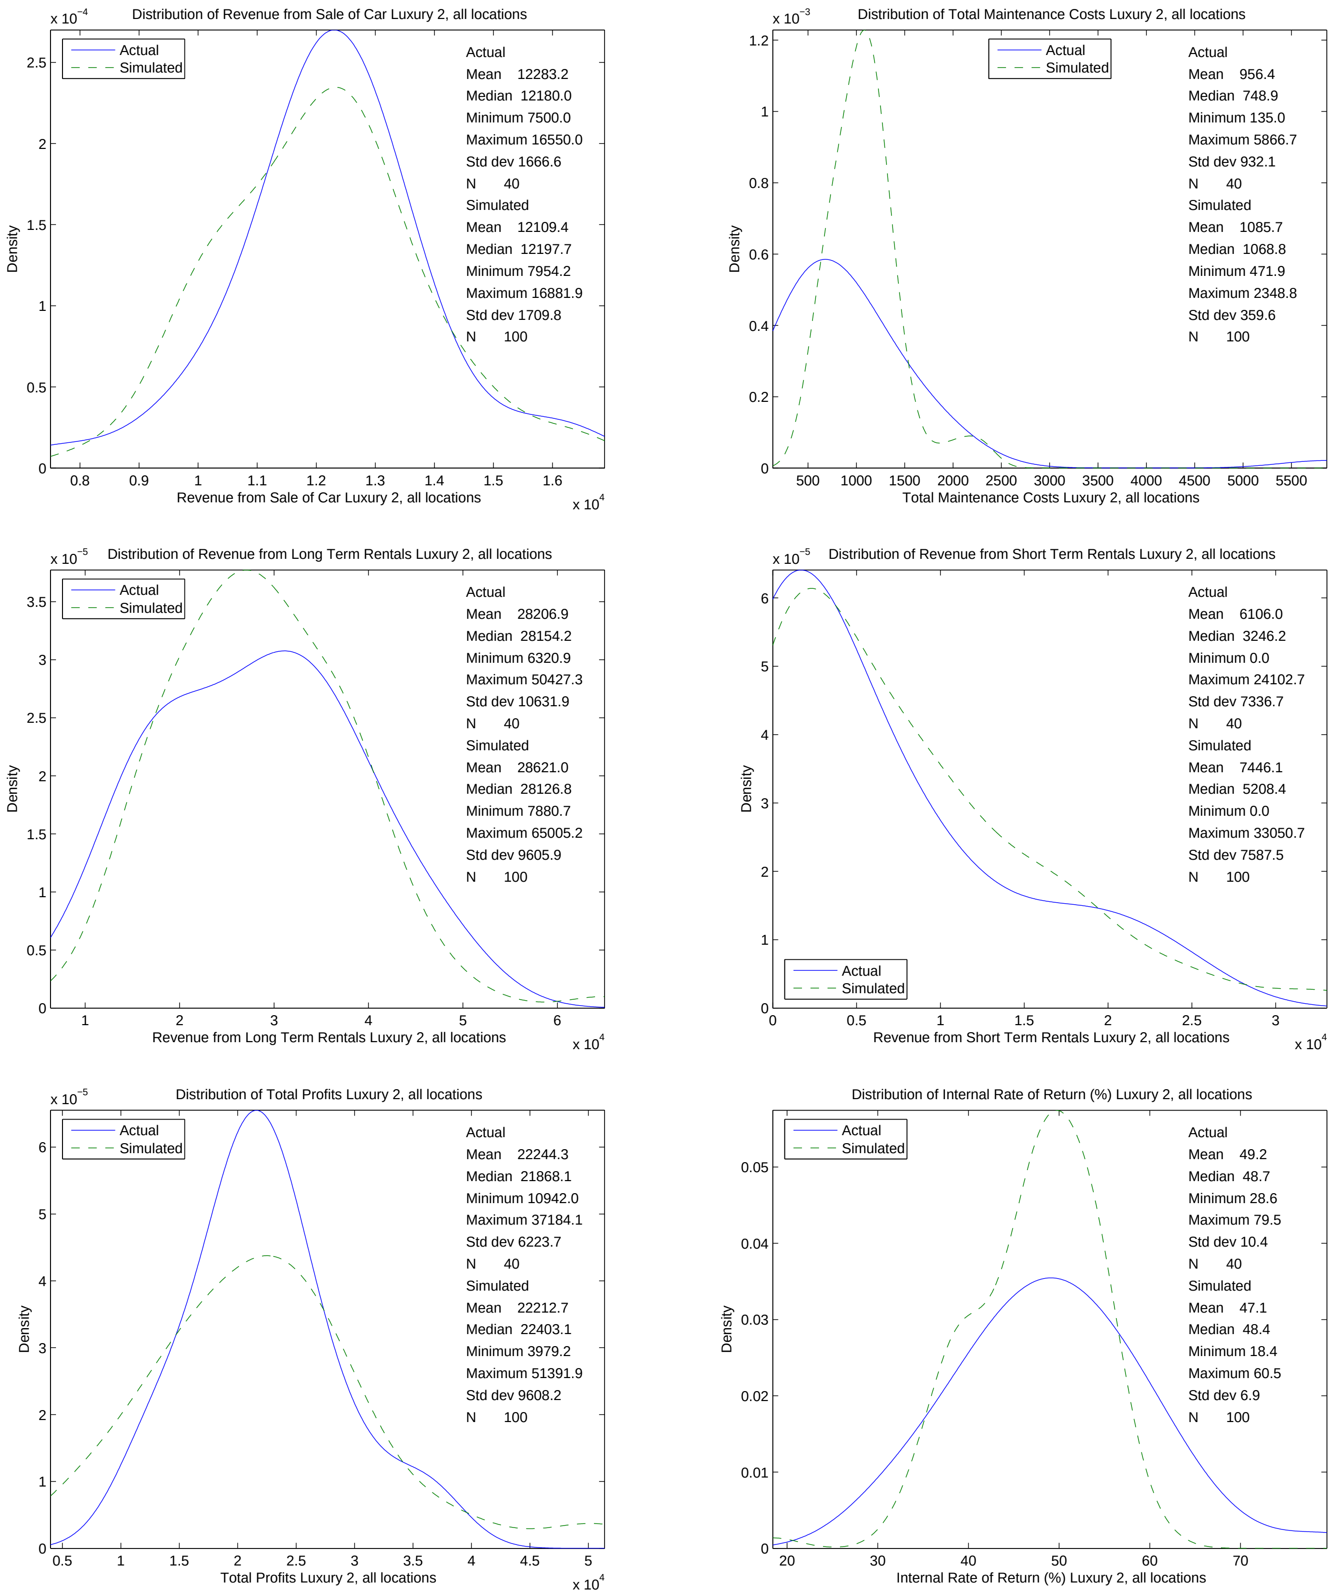

Figure 13 Simulated versus Actual Costs, Revenues and Profits: Luxury – All Locations

nearly 90% of the rental days were in long term contracts. However it might seem surprising in view of our finding that there is a negative relationship between the fraction of time spent in long term contracts and the IRR on the vehicle. This suggests that long term contracts are less profitable than short term contracts, and thus the company should try to allocate a greater share of the vehicle's rentals to short term instead of long term contracts. However as we noted above, there are reasons to distrust any conclusions based on the simple IRR regression in table 1 and we will return to the issue of the relative profitability of long verus short term rental contracts in the next section.

The bottom two panels of figure 13 compare the actual and simulated distributions of total profits and internal rates of return, IRRs. Once again our econometric model provides a very good estimate of mean total profits and mean IRR, although the simulation results result in a distribution of total profits that has a larger variance than the actual distribution, and a distribution of IRRs that 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, but the reassuring thing to notice is that it does correctly predict mean total profits and the mean IRR. Since we model the company as a an expected profit maximizer (i.e. the company is not "risk averse"), only the mean values of profits matters: 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.

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 Characterizing the Optimal Replacement Policy

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.

We use the method of *dynamic programming* to formulate and solve the optimal stopping problem. We show that the optimal strategy 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 $\overline{o}(d,r,\tau)$ that depends on the current rental state r, the duration in that state d, and the car type τ. Using numerical methods, we solve the dynamic programming problem and calculate the optimal stopping thresholds $\overline{o}(d,r,\tau)$ for the compact, luxury and RV car types and the associated optimal *value functions* $V(r,d,o,\tau)$. This function provides the expected discounted profits (over an infinite horizon) under the optimal replacement policy for a vehicle 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 $\mu(r,d,o,\tau)$. We let $V_\mu(r,d,o,\tau)$ denote the expected discounted profits (again over an infinite horizon) under the alternative replacment 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 increase its expected discounted profits by over 10%.

The optimal stopping problem emerges from the solution to the *Bellman equation,* which is a recursively defined formula the value function V given by

$$V(r,d,o)=\max\left[EP(o)-\overline{P}+\beta EV(r_0,0,0),ER(r,d,o)-EM+\beta EV(r,d,o)\right] \tag{14}$$

where we suppress the τ notation under the understanding that separate Bellman equations are solved for each of the three car types $\tau\in\{\text{compact},\text{luxury},\text{RV}\}$. In the Bellman equation (14), $\overline{P}$ denotes the cost of a new car, $EP(o)$ is the expected resale value of a car with odometer value o, $ER(r,d,o)$ is the expected rental revenue from renting the car to customers, EM is the expected daily cost of maintaining the car including the cost of cleaning cars at the end of rental contracts, and EV is the expected discounted value of future profits from operating a sequence of rental cars (possibly until the infinite future).

There are two EV functions in the Bellman equation, $EV(r,d,o)$ and $EV(r_0,0,0)$. The term $EV(r,d,o)$ denotes the expected value of an *existing* car which has an odometer value of o and has been in rental state r for a duration of d days. The term $EV(r_0,0,0)$ denotes the expected value of a *new car* just after it has been purchased when it is on the lot waiting for its first rental. The notation r_0 denotes the first lot spell. To economize on states, we actually assuming that this function can be represented in terms of lot states $r=3$ and $r=4$ (where recall these are lot states where the previous rental spell was either a long term contract or a short term contract, respectively), as

$$EV(r_0,0,0)=[\eta EV(3,0,0)+(1-\eta)EV(4,0,0)] \tag{15}$$

where the parameter $\eta\in(0,1)$ is chosen so that the weighted average duration distributions and the transition probability for the initial lot spell matches the mean duration for initial lot spells that we observe in the data, and results in the same fraction of cars whose first rental spell is a long term rental as we observe in the rental data set.

The left hand term on the right hand side of (14) is the expected value of replacing a current vehicle with a new one. Thus, $EP(o)-\overline{P}$ is the expected cost of replacement, i.e. the expected resale value of the existing car (which has odometer value o) less the cost of a new replacement car $\overline{P}$, plus the expected discounted value from tomorrow onward, $\beta EV(r_0,0,0)$. We assume that a brand new car has an odometer value of $o=0$ and starts its life a lot spell with a duration of $d=0$.

The Bellman equation (14) actually applies only when the car is in a lot spell ($r\in\{3,4\}$) or before the first day of a rental spell ($d=0$ if $r\in\{1,2\}$), since we assume that the company will not interrupt an ongoing rental contract to sell a vehicle. Thus, for cars in the midst of a rental spell, ($r=1$ or $r=2$ and $d>0$), we have

$$V(r,d,o)=ER(r,d,o)-EM+\beta EV(r,d,o). \tag{16}$$

The $EV(r,d,o)$ function is a conditional expectation of the value function $V(r,d,o)$. For lot spells, $r > 2$, we have

$$EV(r,d,o) = h(d,r)\left[V(1,1,o)\pi(1|r,d,o) + V(2,1,o)[1-\pi(1|r,d,o)]\right] + [1-h(d,r)]V(r,\min(d+1,31),o). \tag{17}$$

What this equation says is that with probability $h(d,r)$ the lot spell ends and the car will transit either to a rental spell under a long term contract $r = 1$ with probability $\pi(1|r,d,o)$ or a rental spell under a short term contract with probability $1-\pi(1|r,d,o) = \pi(2|r,d,o)$ since we have ruled out self-transitions back to the lot, $\pi(r|r,d,o) = 0$ for $r > 2$. With probability $1-h(d,r)$, the lot spell continues and the value function in this case will be $V(r,d+1,o)$, reflecting an increment of 1 more day to the duration counter, unless $d \geq 31$ in which case d remains at the absorbing state value of $d = 31$. We accomodate the absorbing state assumption for duration in lot spells by writing $d_{t+1} = \min(d_t+1,31)$ in the equation for the value function above.

For rental spells, $r \in \{1,2\}$, we have

$$EV(r,d,o) = h(d,r)\int_{o'}\left[\sum_{r'} V(r',0,o')\pi(r'|r,d,o)\right] f(o'|r,d,o)do' + [1-h(d,r)]V(l(r),d+1,o), \tag{18}$$

where $l(r)$ is the lot state following a termination of rental state r, i.e. $l(1) = 3$ and $l(2) = 4$, and $f(o'|r,d,o)$ is the conditional density of the number of kilometers on the odometer of a car returning from a rental spell of type r that has lasted d days and started with an odometer value of o. Thus, if a car is in a rental spell, it will either remain in the rental spell for another day with probability $1-h(d,r)$ (unless $d \geq 31$, in which case $h(d,r) = 1$), or with probability $h(d,r)$ the rental spell ends and the car transits to a new rental state r', which will either be a lot spell, $r' > 2$, or a rental spell under a short contract $r' = 2$, or a long term contract $r' = 1$. If a car remains in a rental spell, the company will not know the odometer reading until the car returns from the spell. Thus, we keep the odometer state variable o fixed at its original value as long as a car continues its current rental spell. However if a car returns from a rental spell, the company learns the mileage travelled by the customer, $o'-o$, and thus the odometer state variable increases from o to o'. As noted above, we assume that the total mileage driven under a rental contract that has lasted d days and is of type $r \in \{1,2\}$ is an Erlang distribution with parameters d and λ_r, where λ_r is the mean mileage driven per day by customers under contract type r.

Finally we specify the expected rental revenue function, $ER(r,d,o)$. Initially we assume that long term and short term contracts allow unlimited kilometers and are charged on a daily rate, except with an early return penalty for long term contracts. Thus for short term contract, $r = 2$, we have

$$ER(2,d,o) = h(d,2)EDR(2)d, \tag{19}$$

where $EDR(2)$ is the expected daily rental rate for a short term contract. We multiply by the hazard since we assume that the rental is paid only at the end of the rental spell, but no revenue is received otherwise (i.e. if the rental contract continues another day). This expected value accounts for cases where cars are rented multiple times in the same day as "chauffered vehicles" and reflects the expected sum of all rental revenue received on such days less the amount paid to the chauffer.

For long term rentals, there is a lower per day rate, $EDR(1)$ provided the vehicle is rented for sufficiently many days, say $\overline{d}$. Otherwise the car is treated as an early return and there is a per day penalty rate, ρ, added on for such early returns. Thus, the expected revenue function is

$$ER(1,d,o) = \begin{cases} h(d,1)EDR(1)d & \text{if } d \geq \overline{d} \\ h(d,1)[EDR(1)+\rho]d & \text{if } d < \overline{d} \end{cases} \tag{20}$$

Now consider the non-regenerative optimal stopping formulation of the car replacement problem. This formulation of the problem differs from the regenerative optimal stopping by focusing only on the sale of the existing car: it ignores the issue of replacing the current car with another new one when the current is sold, and the infinite sequence of replacements over the infinite future. In this formulation, the problem is to determine a stopping time (i.e. time to replace the car) that maximizes the firm's expected discounted profit from operating that particular car, but we do not assume that when this car is sold that it will be replaced by a new one. This effectively makes the problem a finite horizon problem, but with a random horizon (the horizon equals the random stopping time). The Bellman equation for this problem is

$$V(r,d,o) = \max\left[\overline{E}P(o), ER(r,d,o) - EM + \beta EV(r,d,o)\right] \tag{21}$$

Comparing this to the Bellman equation for the infinite horizon regenerative optimal stopping formulation of the problem we see that the value of stopping is equal to $EP(o)$, the expected selling price of the car, whereas in the regenerative optimal stopping problem the value of stopping equals $EP(o) - \overline{P} + \beta EV(r_0,0,0)$ which equals the expected selling price plus the "net continuation value" $\beta EV(r_0,0,0) - \overline{P}$. The latter value equals the expected discounted value of replacing the car with a new one, less the cost of that car. The continuation value actually represents the net value of an infinite stream of future replacements (at random future stopping times), and is generally expected to be positive.

While it is possible that more complicated types of optimal replacement rules could arise, the optimal replacement policy will generally takes the form of an *optimal stopping threshold* $\overline{o}(r,d)$. This is simply a rule that says that it is optimal to keep the current car if its odometer o satisfies $o < \overline{o}(r,d)$ and to replace the car otherwise. The optimal stopping threshold is the value of o where the firm is indifferent between keeping the car and replacing it. That is, it is the solution to the equation

$$EP(\overline{o}(r,d)) - \overline{P} + \beta EV(r_0,0,0) = ER(r,d,\overline{o}(r,d)) - EM + \beta EV(r,d,\overline{o}(r,d)), \tag{22}$$

in the regenerative optimal stopping formulation of the optimal replacement problem and

$$EP(\overline{o}(r,d)) = ER(r,d,o(r,d)) - EM + \beta EV(r,d,o(r,d)), \tag{23}$$

Note that $EP(o)$ is a decreasing function of o, but $\overline{P} - \beta EV(r_0,0,0) + ER(r,d,o) - EM + \beta EV(r,d,o)$ will also be a decreasing function of odometer, o. However in this problem the latter function decreases more rapidly as a function of o than $EP(o)$ and there is generally a unique solution $o(r,d)$ to both equations above.

Since $-\overline{P} + \beta EV(r_0,0,0) > 0$ (this condition tells us that the present value of future operating profits exceeds the purchase price of a car, otherwise if this condition is not satisfied, the company would not replace the current car once they sell it), it would appear that $\overline{o}(r,d)$ should be *lower* in the regenerative optimal stopping formulation of the replacement problem (22) since the left hand side of this equation is a parallel upward shift of the $EP(o)$ function, which is less steeply sloped (in o) than the right hand side terms, $ER(r,d,o) - EM + \beta EV(r,d,o)$.

However there is a complication with this reasoning since the EV functions on the right hand sides of equations (22) and (23) are not the same: there are two different value functions that we are taking expectations of. To be clearer, let the expected value in the non-regenerative optimal stopping problem be $EV_{fh}(r,d,o)$ and the expected value in the regenerative optimal stopping problem be $EV_{ih}(r,d,o)$, where ("ih" denotes infinite horizon and "fh" denotes finite horizon). It is not difficult to show that $EV_{ih}(r,d,o) > EV_{fh}(r,d,o)$, so whether $\overline{o}_{ih}(r,d) \geq \overline{o}_{fh}(r,d)$ is theoretically ambiguous: it depends on whether the upward shift in left hand side of equation (22) relative to equation (23) is larger than the upward shift in the right hand side of equation (22) relative to equation (23).

We will show in the next section that in fact, the optimal stopping threshold for the regenerative optimal stopping formulation of the optimal replacement problem is indeed lower than the optimal stopping threshold for the non-regenerative optimal stopping formulation of the problem. The intuition for the result is that delaying replacement a car is more costly for the firm in the regenerative optimal stopping problem, since it postpones the net cash flows arising from purchases of all subsequent replacement cars. Thus, there is an added *opportunity cost* from delaying replacement of a vehicle in the regenerative optimal stopping formulation of the problem compared to the non-regenerative optimal stopping formulation, and this opportunity cost makes it optimal to replace cars *sooner:* i.e. $\overline{o}_{ih}(r,d) < \overline{o}_{fh}(r,d)$.

We believe that the regenerative optimal stopping formulation of the optimal replacement problem is the appropriate version to consider, since in fact, this firm is a continuing business that replaces the vehicles in its fleet as opposed to just selling them off and not replacing them. However the non-regenerative stopping formulation is relevant to describe optimal behavior in a *shut down scenario,* i.e. when $\overline{P} > \beta EV(r_0,0,0)$, the firm will want to operate its existing rental vehicles optimally over their remaining lifetimes, but without replacing them at the time they are sold. In the next section we will show that $\beta EV(r_0,0,0)$ is many times larger than $\overline{P}$, so in fact, the shut down scenario is not relevant for any of the car types we have analyzed.

6 Calculating the Optimal Replacement Policy: Numerical Results

As we noted above, if we 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" (see section 4), then the optimal stopping thresholds is $\overline{o}(r,d) = \infty$, i.e. it is *never optimal to sell an existing vehicle.* These results are due to our 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* (see the discussion in section 5), 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 our data.* That is, we will assume that beyond the range of our 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. Figure 14 displays the assumptions we have used. The left hand panel of figure 14 shows the multiplication factor that we apply to daily maintenance costs as a function of the vehicle's odometer. Thus, consistent with our data, over the range from $o \in [0,130,000)$ kilometers, we assume daily maintenance costs do not increase, however outside the range of our data, we

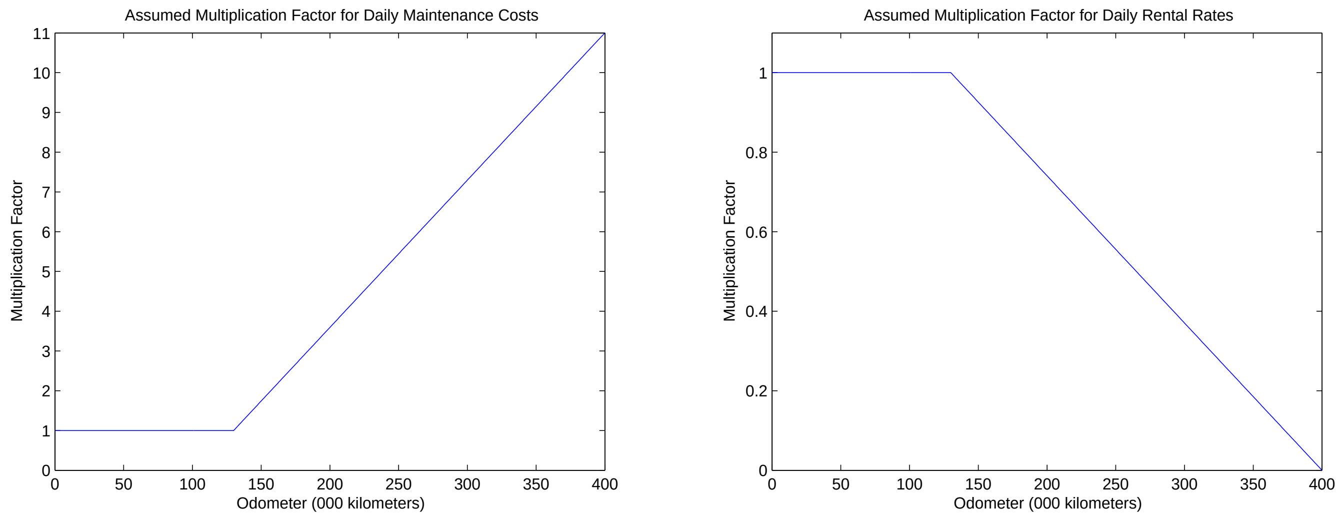

Figure 14 Multiplication Factors for Daily Maintenance Costs and Rental Rates: All Car Types – All Locations

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. Indeed, as we noted in the introduction, the company conducted an experiment to discount rental rates for vehicles over 2 years old, but stopped it when it found that virtually all customers preferred to rent an older car at a 20% discount rather than a newer car at the full daily rate. This suggests that discounts could be popular with many customers and that required discounts would be much less drastic than what we have assumed.

Figure 15 shows our calculated optimal replacement thresholds $\overline{o}(d, r, \tau)$ and the value functions $V(d, r, o, \tau)$ for the three car types $\tau \in \{\text{compact}, \text{luxury}, \text{RV}\}$. The left hand panels display the replacement thresholds

$\overline{o}(d,r,\tau)$. The solid blue line is the replacement threshold for the case where a car is in a lot spell of type $r = 3$, i.e. the previous rental spell was a long term contract, and the red dashed line is for a car that is in a lot spell of type $r = 4$, i.e. where the previous rental was a short term contract. In addition, the firm can decide whether or not to replace a car at the start of a rental spell, and the figures also present the these thresholds, $\overline{o}(0,1,\tau)$ (the threshold applicable at the start of a long term rental spell) and $\overline{o}(0,2,\tau)$ (the threshold applicable at the start of a short term rental spell).

We see that for type $r = 3$ lot spells, the replacement threshold is approximately equal to 145,000 kilometers for all three car types. This threshold is basically flat as a function of the duration in the spell, except that in the case of the RV, the threshold starts out at about 157,000 kilometers and then decreases with the duration on the lot to about 146,000 kilometers for cars that have been on the lot 30 days or more. The fact that the replacement threshold is essentially flat as a function of duration in the lot for a type $r = 3$ lot spell is due to two factors: 1) the hazard rate for exiting the lot is essentially flat as a function of duration on the lot d, and 2) there is a very high probability that at the end of the lot spell the car will transit to another long term rental spell. These two factors imply an absence of "duration dependence" (i.e. length of time on the lot does affect the chances that the vehicle will leave the lot the next day) which in term implies an absence of duration dependence in the replacement threshold $\overline{o}(d,3,\tau)$.

However for type $r = 4$ lot spells, i.e. lot spells that were preceded by a short term rental spell, there is duration dependence: i.e. the probability of leaving the lot spell is a decreasing function of the length of time spent on the lot d. In addition, there is a significantly lower chance that the next rental spell will be a long term rental, and thus a high chance that the vehicle will have a relatively rental spell in a short term contract and will return to the lot again. Thus, a vehicle that is current in a type $r = 4$ lot spell is more likely to be idle in the future, and this fact, combined with the decreasing chance of exiting the lot as lot duration d increases, causes the replacement threshold $\overline{o}(d,4,\tau)$ to decline with d. For all three car types, we see fairly steep declines in $\overline{o}(d,4,\tau)$. For example, for the compact, the replacement threshold starts out at approximately $164,000$ kilometers on the first day of the spell, and declines to about $144,000$ kilometers by the 31st day on the lot.[7]

It is also not surprising that the replacement thresholds are significantly higher at the start of a short or long term rental spell. The company has the option to sell the car at this point, or to let the renter take it. It seems intuitively obvious that the replacement threshold should be higher when the company has "a bird in

[7] Due to our constant hazard assumption for lot spell durations longer than $d = 31$, this $144,000$ kilometer threshold applies to all durations over $d = 31$.

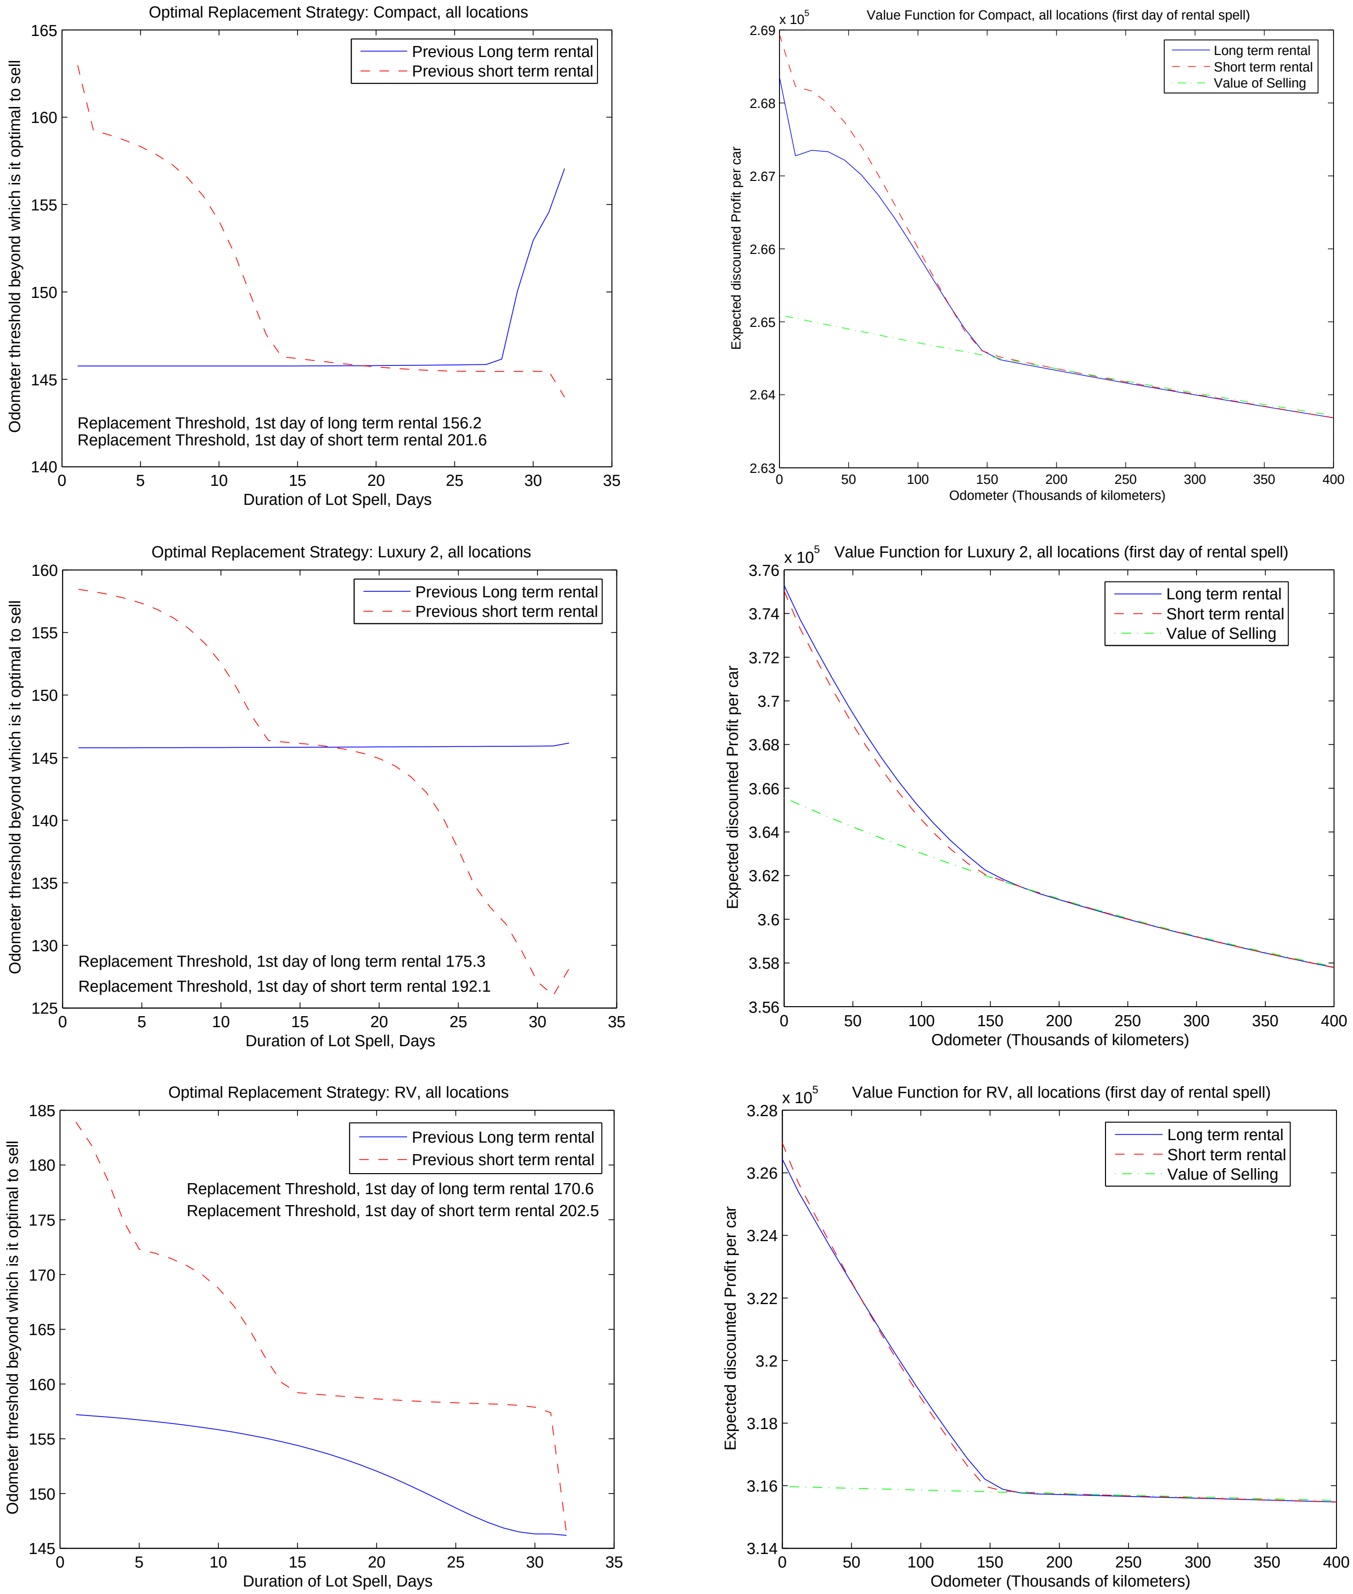

Figure 15 Optimal Decision Rules and Value Functions: All Car Types – All Locations

the hand" than for a car that is on the lot waiting to be rented. For the compact, the replacement threshold at the start of a long term rental contract is $\overline{o}(0,1,\tau) = 156,200$, which exceeds the $145,000$ kilometer threshold for a car in the first day of a lot spell that has just emerged exited a long term rental spell. We see that the replacement threshold at the start of a short term rental contract is even higher. In the case of the compact the threshold is $\overline{o}(0,2,\tau) = 201,600$ which is significantly higher than the $164,000$ kilometer threshold for a car that is fresh on the lot, having just emerged from a short term rental spell. The disparity in replacement thresholds at the start of long and short term rental spells is not as great for the other two car types. The explanation is that short term contracts are much more lucrative than long term contracts for the compact car type compared to the other two car types, since the ratio of daily rental rates for short term contracts to long term contracts is significantly higher for the compact car type than for the other two car types. Thus, the company is relatively more eager to take advantage of a short term rental opportunity for its compact cars, and this makes it optimal for it to have a significantly higher replacement threshold at the start of a short term rental contract.

The right hand panels of figure 15 plot the calculated value functions for each of the three car types as a function of odometer o. We plot three functions: the value of keeping a car on the first day of a rental spell, $V(o,1,r,\tau)$ where $r = \{3,4\}$, and the expected value of selling the vehicle, $EP(o,\tau) - \overline{P}(\tau) + \beta EV(0,r_0,0)$ as a function of its current odometer value o. Clearly, if the value of keeping the car exceeds the value of selling it, then it is optimal to keep it. Thus, the first point at where the value of keeping the car equals the value of selling it constitutes the optimal replacement threshold $\overline{o}(d,r,\tau)$. Thus, for the compact, if we were to enlarge the top right hand panel of figure 15 we would see the red dashed line (the value of keeping a car when it is in lot spell type $r = 4$) first intersects the green dash-dotted line at 201,600 kilometers, so this constitutes the optimal replacement threshold that is appropriate at the start of a short term rental spell that we noted previously. The solid blue line is the value of keeping a car when it is in lot spell type $r = 3$, and it first intersects the green dash-dotted line at 156,200 kilometers, which is the optimal replacement threshold at the start of a long term rental spell.

For all three car types, the "rental value", i.e. the difference in the value of keeping and selling it for a new one is a relatively steeply decreasing function of odometer value. While we noted that the resale value of a car is a sharply decreasing function of the vehicle's odometer (except in the case of the RV where the rate of depreciation is much flatter after the initial rapid depreciation that occurs in the first 20,000 kilometers), *we see 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 our econometric analysis in section 3. Note that our assumed sharp drop off in rental rates and sharp rise in maintenance costs that we displayed in figure 14 do not start until after 130,000 kilometers, yet the decline in value of a car occurs immediately. The only aging effects we uncovered in our 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 values of rental vehicles that we see in figure 15.

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 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. Thus, the values of keeping an existing car in figure 15 represent *expected future profits* over the life of the car, but the new purchase price of the *current* vehicle is treated as a a “sunk” or “bygone expense.” Thus, as the vehicle’s odometer increases from zero towards the optimal replacement threshold, the expected discounted value of remaining profits *on the current car* necessarily decreases since there is a shorter remaining life *for the current car* over which this profit stream will be collected. Then when the company finally replaces the vehicle (the value of which is represented by the green dash-dotted lines in figure 15), the firm does actually incur the cost of buying a new replacement vehicle and the process starts over again. Thus, the difference between the value of keeping a (just purchased) new vehicle and immediately trading it for another new vehicle, $V(0,r_0,0) - [EP(0) - \overline{P} + \beta EV(0,r_0,0)]$, represents the expected discounted profit that the firm expects to earn on the current vehicle over its lifetime. Thus, for the luxury car type, we have $V(0,r_0,0) = 375,000$ whereas the value of immediately selling a newly purchased car is approximately $EP(0) - \overline{P} + EV(0,r_0,0) = 366,000$. 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.

Figure 16 plots value functions and expected revenues as a function of duration of the spell, d, keeping the other state variables r and o fixed. We fix o at 47,500 kilometers and plot 4 lines for $V(d,r,o)$ as a function of d for the 4 possible values of $r \in \{1,2,3,4\}$. We see that for the compact car type, the highest value function is for when the car is in a short term rental spell $r = 2$, depicted by the dashed red line in the top left panel of figure 16. This is significantly higher than the value of being in a long term rental spell, which is depicted by the solid blue line. On the other hand, for the luxury car type, the middle left hand panel of figure 16 shows that the long term contracts are more profitable than short term contracts. The RV is similar to the compact in the sense that short term rental contracts are also more profitable than long term contracts.

The right hand panels of figure 16 provides insight into why short term contracts are relatively more profitable than long term contracts for the compact and RV car types, but not for the luxury. The right hand panels plot the expected rental revenues net of expected daily maintenance costs as a function of the duration of the rental spell. Recall that the firm does not actually collect the rental revenues until the end of the rental spell, at which time the total duration of the rental spell becomes known. Thus, the *expected* rental revenues are given by formulas (19) and (20) in the previous section, as a product of the daily rental rate, times the duration *d times the hazard function* $h(d,r)$. The hazard function is required because it provides the probability that the rental spell ends on the $(d+1)^{\text{st}}$ day, so that with probability $h(d,r)$ the company collects total revenue $ER(r)d$ and with probability $1-h(d,r)$ it collects zero (since the latter case provides the probability that the car does not return on day $d+1$, so no rental revenues would be collected yet).

We see that expected rental revenues for a long term rental spell, depicted by the solid blue line, are essentially zero until around 28 days, when the expected revenues start to rise. The reason that expected revenues are zero prior to 28 days can be easily seen from the solid blue line in the left hand panel of figure 5: this is the duration distribution for rental spells and it shows that there is essentially zero probability that a car in a long term rental will be returned prior to 28 days. By the longest possible duration, $d = 31$, the hazard rates rise to 1, and so the expected rental revenues increase monotonically to the total revenues for the maximum rental term, $31 * ER(r)$ for $r = 1$ and $r = 2$. The ratio of these "full term rental revenues" is simply equal to the ratio of the daily rental rates, and since the company charges a higher per day rate for short term rentals than for long term rentals, the revenues for short term contract that lasts 31 days are always significantly higher than for a long term contract that lasts 31 days.

Of course, it is very unlikely that a short term contract will last 31 days: as you can see from the left hand panel of figure 5, the duration distribution for short term rentals has most of its mass for durations of 5 or fewer days. Thus, the *ex ante expected duration* (i.e. not conditioning on duration, d) of a short term rental spell is only a few days, whereas the expected duration of a long term rental contract is about 30 days. Thus, even though short term contracts have higher daily rental rates, it is not clear *a priori* whether they are more profitable than long term contracts: we need to account for a greater probability of idle

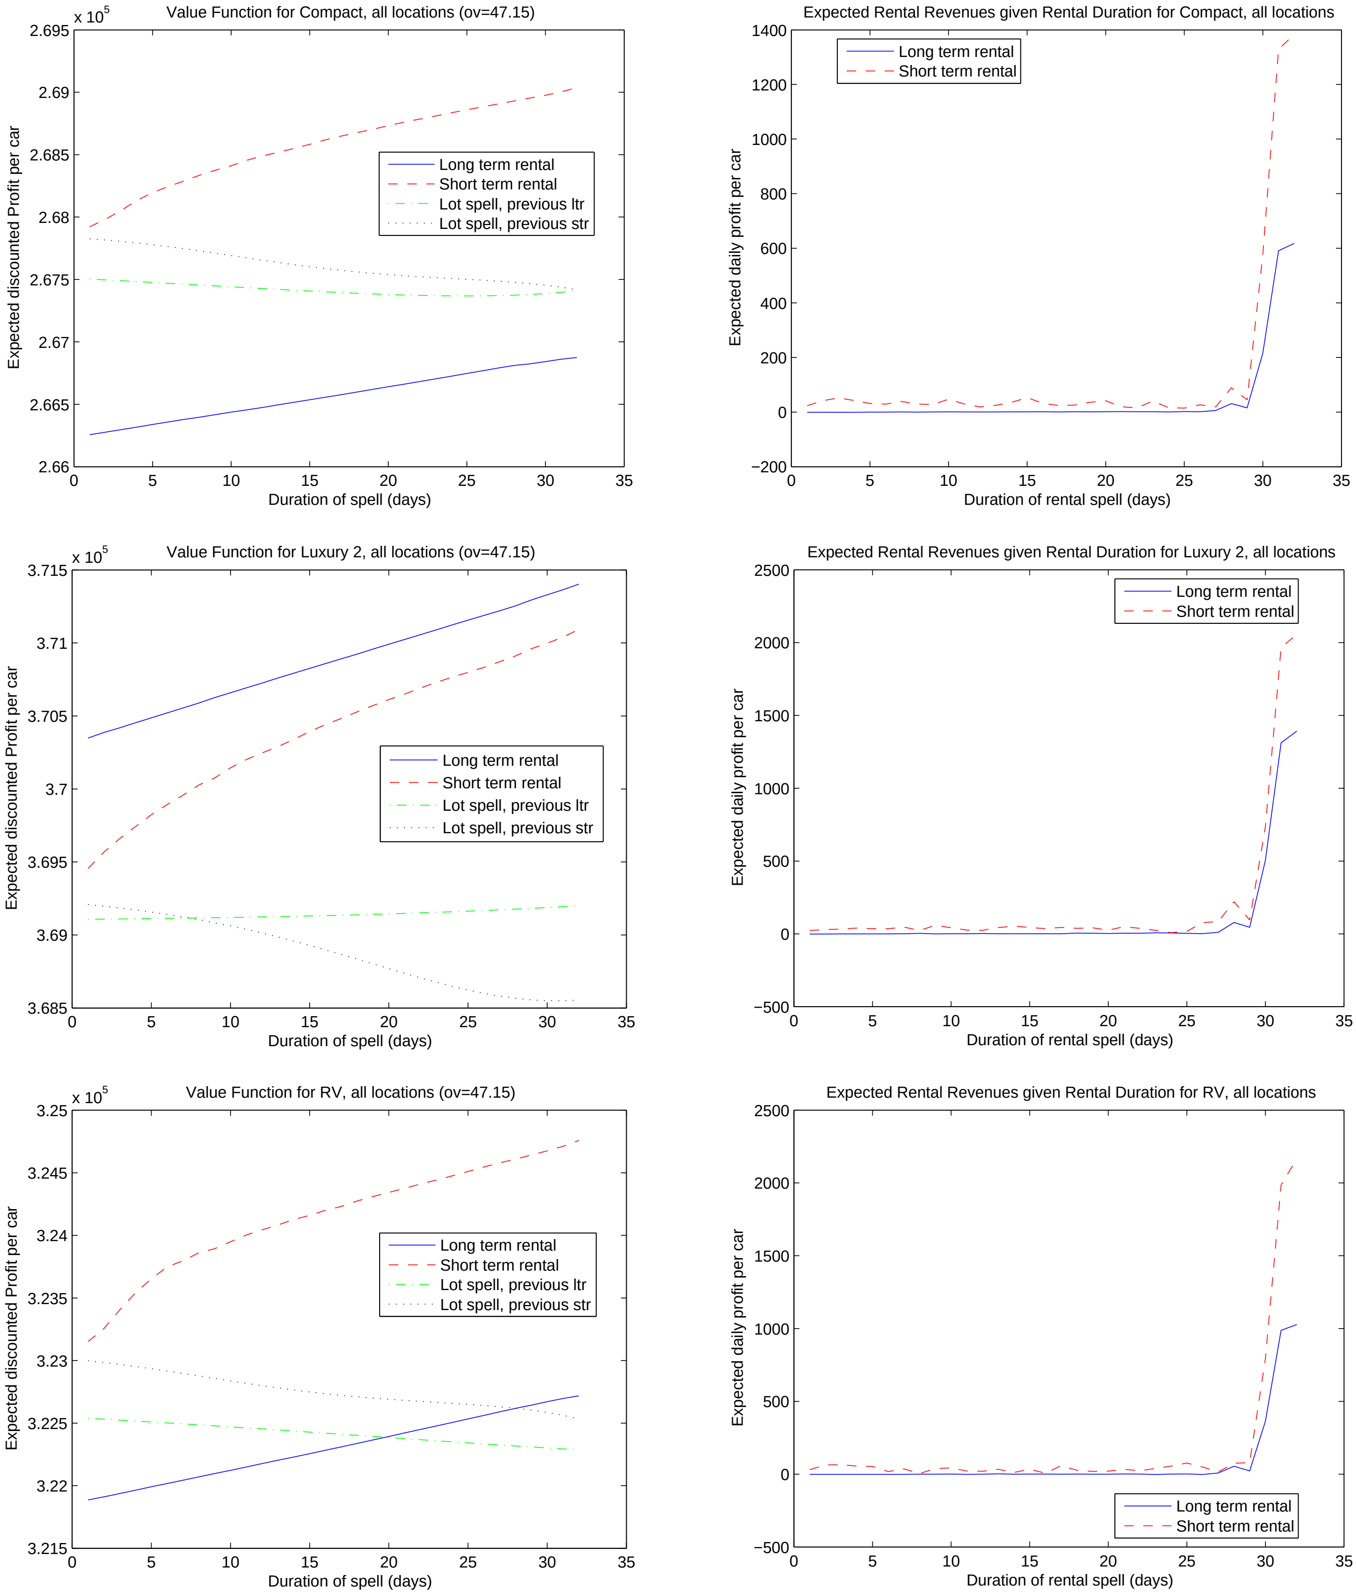

Figure 16 Expected Rental Revenues and Value Functions: All Car Types – All Locations

periods between successive short term contracts when comparing their profitability to long term contracts. Recall that our regression analysis results in table 1 predicted that long term contracts where uniformly less profitable, since the regression coefficient on the variable “Fraction Rented Long Term” was negative (although we note that it was sigificantly negative for the compact and RV car types, but not significantly different from zero in the case of the luxury car).

However we can see from figure 16 that at least under an optimal replacement policy, the long term contract is predicted to be more profitable for the luxury, but less profitable for the compact and RV. This finding is an interaction of the relative rates the company charges for short term versus long term rentals: for luxury vehicles, the company’s rental rate for long term contracts is 67% of the daily rate it charges for short term rentals, whereas for the compact and RV, the company’s daily rate for long term contracts is only 45% and 47% of the corresponding short term rates, respectively. Thus, the company’s rental rate structure discounts rental rates for long term rentals of luxury cars less than it does for compact or RVs. This is part of the explanation why long term rentals are more profitable for the luxury car type. The other part of the explanation is more subtle, but it lies in the estimated transition probabilities and hazard rates: there is a higher probability of “roll over” of long term rental contracts for the luxury compared to the other two car types, and a higher probability that a car in a lot spell will return to another long term contract for luxury car types than for compacts or RVs. Thus, long term rentals of the luxury car type have a higher “effective capacity utilization rate” per unit time than for short term rentals, so this is the other part of the explanation why long term rentals are more profitable for luxury cars, but not for compact or RVs.

The other thing to notice about the left hand panels of figure 16 is that we generally find the expected result that the value of being in a lot spell is lower than being in a rental spell of the corresponding type (i.e. being a long term rental spell is more profitable than being in lot spell where the previous rental was long term, and being in a short term rental spell is more profitable than being in a lot spell when the previous rental spell was short term). However, we see in the top left hand panel of figure 16 that this does not hold in the case of a compact: the value of being in a type $r = 3$ lot state is *higher* than the value of being in a type $r = 1$ rental spell! This is due to the fact that long term rental contracts are so much less profitable than short term rental contracts for the compact car types. When compact car is in a lot spell, there is a better chance that it will make a quick transition to a much more profitable short term rental spell than if it is in a long term rental contract. We also see that the value of being in a type $r = 4$ lot spell (i.e. where the previous rental spell was a short term contract) is higher than being in a type $r = 3$ lot spell. This is because a type $r = 4$ lot spell has a higher probability of transition into to a short term rental spell than a type $r = 3$ lot spell.

In order to learn more about the implications of the optimal replacement policy, particularly about the distribution of ages at which replacements occur, we resorted to stochastic simulation of the optimal replacement policy using the same basic approach that we used to simulate the econometric model under the *status quo* replacement policy in section 4. Figure 17 presents a comparison of the simulation results (for 100 simulated cars of each of the three car types) and the firm's actual outcomes under its *status quo* replacement policy. The left hand panels compare the simulated and actual distribution of odometer values of cars at replacement, and the right hand panel compares the distributions of ages at replacement.

We see 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 three right hand panels of figure 17 compare the actual distribution of ages at replacement to the distribution predicted to occur under the optimal replacement policy. We see that under the optimal replacement policy the mean age at replacement of all three car types is almost twice as high: it 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. We can see from the right hand panels of figure 17 that 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 we observe for vehicles we observe that have fewer than 130,000 kilometers. So the combination of the rental discounts and rapid increase in maintenance cost do take its toll, and greatly alter the optimal replacement policy. Instead of it being optimal to *never* replace its existing vehicles, once we make the assumptions about rapidly rising maintenance costs and rapidly declining rental rates after 130,000 kilometers, it is no longer optimal for the company to keep and maintain its existing stock of cars indefinitely. However what is surprising is that despite our extremely conservative assumptions, the optimal replacement policy still entails keeping cars about twice as long as the company currently keeps them.

Figure 18 compares the actual distributions of internal rates of returns that the company realizes on

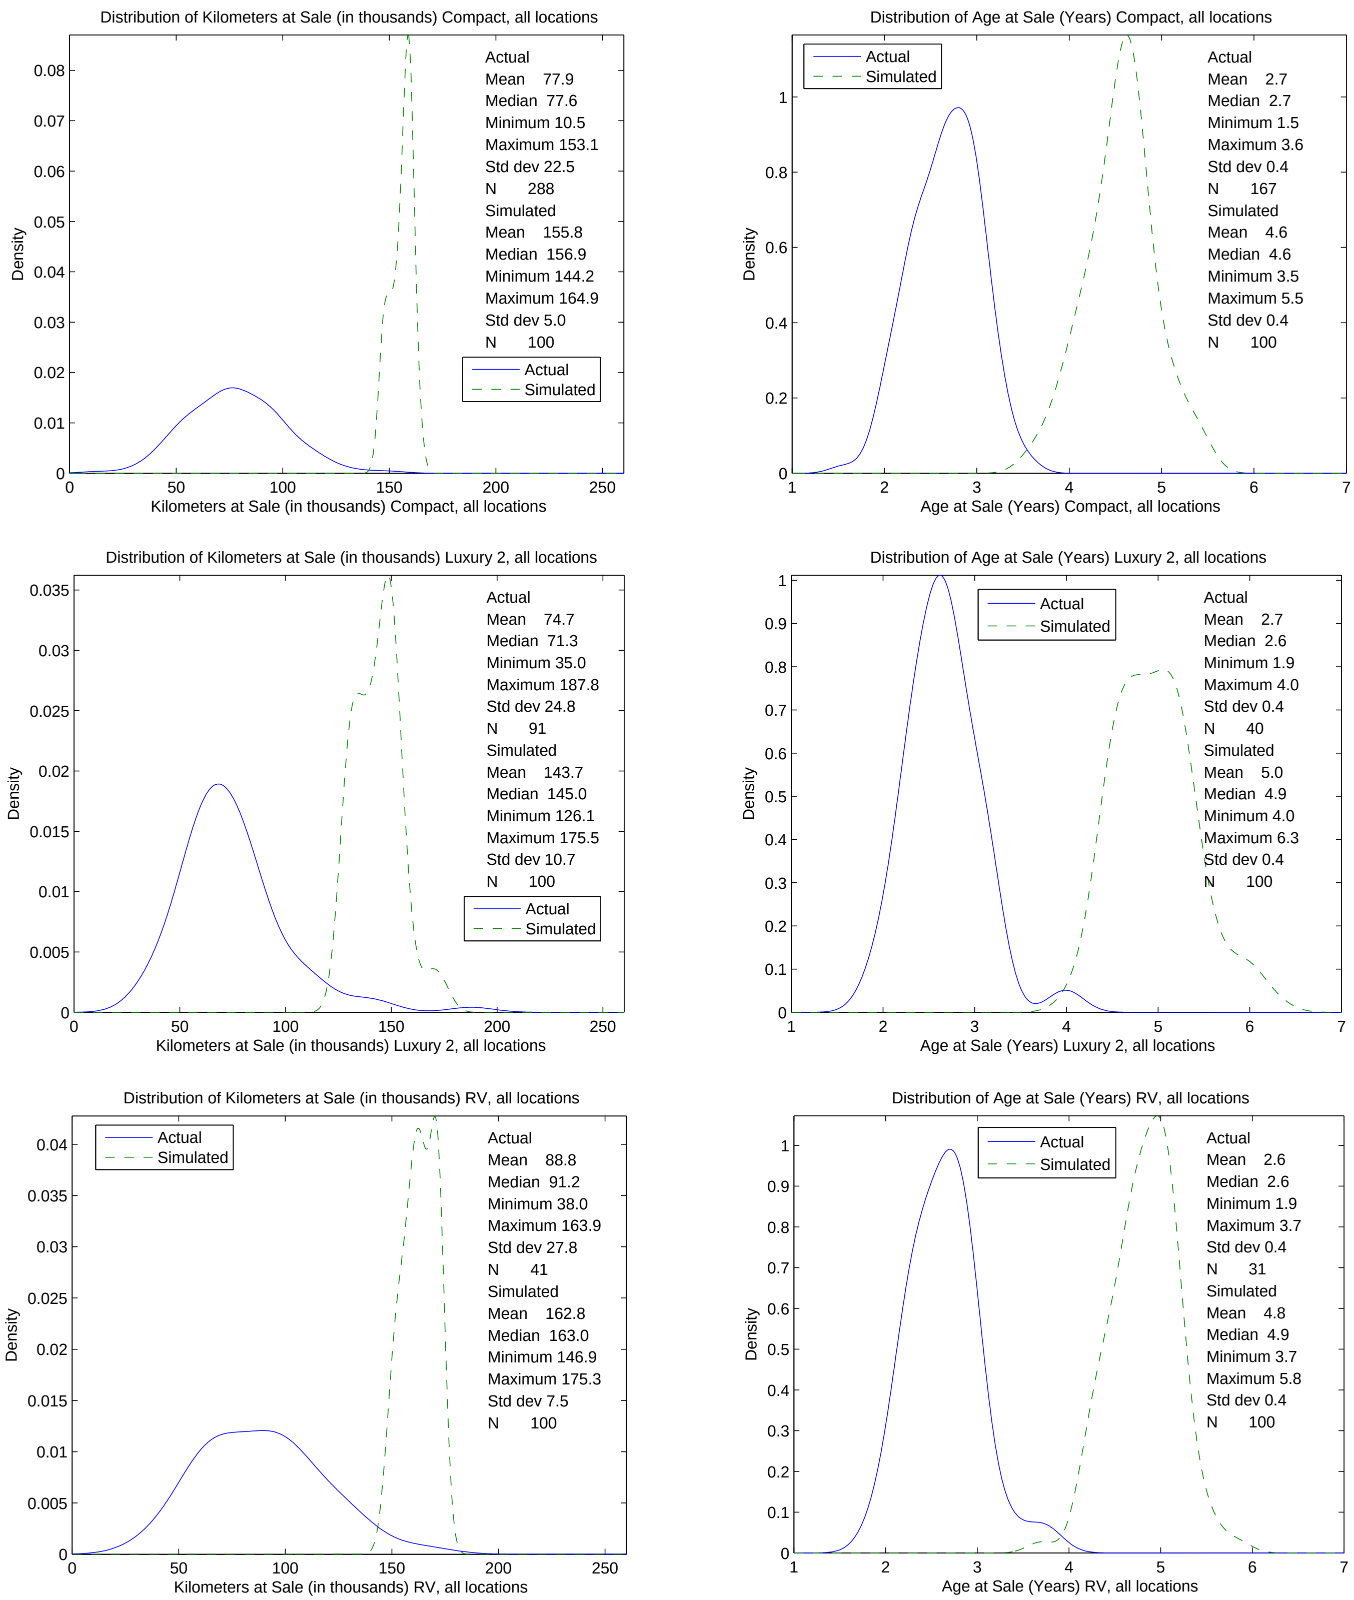

Figure 17 Simulated versus Actual Ages and Odometers at Replacement: All Car Types – All Locations

each of the three car types to the distributions of returns that result under the optimal replacement policy. We see that mean returns are uniformly higher for all three car types. For example, the mean IRR for the compact increases from 77% to 94%, the mean IRR for the luxury increases from 49% to 56% and the mean IRR for the RV increases from 53% to 64%. We see that the distribution of returns has a greater variance under the *status quo* replacement policy, and due to this higher variance, a small fraction of cars, roughly 5%, achieve rates of return under the *status quo* that are higher than the highest possible return earned under the optimal replacement policy. On the other hand, the upper 95% of the distribution of returns under the optimal replacement policy is higher than the median return earned under the *status quo.* Overall, we conclude that adoption of the optimal replacement policy results in a fairly significant upward shift in the distribution of realized returns.

However a more straightforward way to compare the company's *status quo* operating policy with the optimal replacement policy is to compute and compare the *expected discounted profits* under the two operating strategies. We feel that a "fair" comparison requires calculating the expected discounted profits *over an infinite horizon* rather than comparing profits for only a single generation of vehicles. The reason is that comparison that looks only at a single generation will always be biased towards strategies that keep vehicles longer. The reason is, as we discussed previously, the company incurs a large "up front" depreciation cost when it buys a new car. Thus, by keeping vehicles longer the company earns more rental revenues that help to increase profits by "amortising" the initial depreciation expense over a longer service life.

What a "single generation" analysis fails to account for, however, is that by postponing a replacement, *the profits from subsequent generations of vehicles are also postponed, and the postponement of these future profits can represent a large "opportunity cost" that can outweigh the increased short term profits from keeping vehicles longer.* This is exactly the same sort of reasoning underlying our preference for the infinite-horizon regnerative optimal stopping formulation of the optimal replacement problem over the finite-horizon optimal stopping formulation that we discussed in the previous section.

However 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 our data obviously only covers a relatively short time span of the firm's operations. To make this calculation, we solved the following analog of the Bellman equation for the value functions $V_\mu(d,r,o,\tau)$, $\tau \in \{\text{compact}, \text{luxury}, \text{RV}\}$ where we

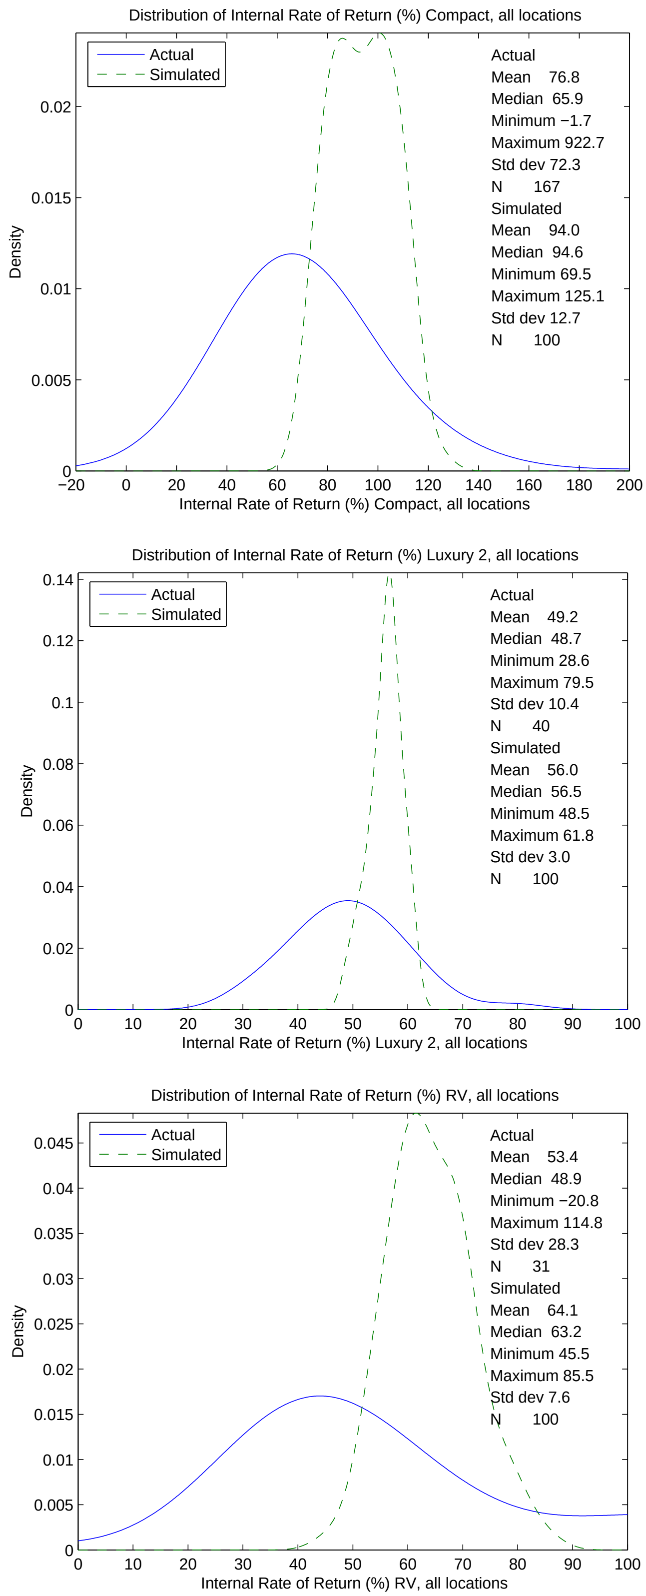

Figure 18 Simulated versus Actual Internal Rates of Return: All Car Types – All Locations

drop the dependence on τ in the equations below to economize on space

$$V_\mu(d,r,o) = \mu(d,r,o)[EP(o) - \overline{P} + \beta EV_\mu(0,r_0,0)] + [1-\mu(d,r,o)][ER(d,r,o) - EM + \beta EV_\mu(d,r,o)], \quad (24)$$

where $\mu(d,r,o)$ is the conditional probability that a car in state (d,r,o) is replaced under the *status quo* replacement policy. We can solve this equation (which is a linear functional equation) using the same numerical techniques that we used to solve for the value function under the optimal replacement policy. By definition, it must be the case that the optimal replacement policy results in higher expected discounted profits than under any alternative strategy μ, so we have

$$V(d,r,o) \geq V_\mu(d,r,o), \quad \forall\,(d,r,o). \quad (25)$$

However the question is "how far is the company's existing replacement policy from optimality?" We can answer this by computing the ratio $V(d,r,o)/V_\mu(d,r,o)$, which represents the factor by which the firm can increase its discounted profits by adopting an optimal replacement policy. If this ratio is not too much bigger than 1, then the company can rest comfortably that its *status quo* replacement policy is "almost" optimal. If the ratio is substantially bigger than 1, then the company might want to re-evaluate its operating strategy.

To simplify our analysis, we focus on comparing the value of a newly purchased brand new car that has just entered the lot. Thus, in table 5 below we report $V(0,r_0,0)$, the value of a new car that has just arrived in the lot under the optimal replacement policy, and $V_\mu(0,r_0,0)$, the value of a new car that has just arrived in the lot under the the firm's *status quo* operating strategy μ. For μ we used the cumulative distribution function for replacements as a function of odometer o only: our approximation to the firm's replacement policy does not depend on spell type r or duration in the spell d. This is justified by our findings in table 2 that once we account for odometer o, neither d and r provide huge enhancements in our ability to predict when the firm replaces one of its vehicles. Table 5 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, $\beta = 0.99991781$. According to the *final value theorem,* (see, e.g. Howard, 1971, p. 46) for any convergent sequence $\{a_t\}$ we have

$$\lim_{\beta \to 1} (1-\beta) \sum_{t=0}^{\infty} a_t = \lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} a_t. \quad (26)$$

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.[8]

The first section of Table 5 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 5 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 our 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 our suggested "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 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
$\overline{P}$	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
$V(0,0,r_0)/\overline{P}$	27.8	16.0	17.4
Expected Discounted Values Under *Status Quo* 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
$V(0,0,r_0)_\mu/\overline{P}$	20.3	13.6	7.3
Ratio of Expected Values: Optimal Policy versus *Status Quo*			
$V(0,0,r_0)/V_\mu(0,0,r_0)$	1.37	1.18	2.39

Table 5: Comparison of Profits/Returns: Optimal Policy versus *Status Quo*

[8]The statement of the stochastic version of the final value theorem is more complex and we omit doing so to save space, but the basic result is the same as the deterministic version of the final value theorem given above.

We conclude this section by undertaking another set of discounted profit calculations to show that our conclusions are robust to even *more conservative assumptions about the increase in maintenance costs and required discounts in rental rates.* Figure 19 presents a second set of more conservative multiplication factors for maintenance and rental rates. Under our second scenario, we assume that maintenance costs begin rising steeply even earlier, at 60,000 kilometers. We emphasize that this assumption is inconsistent with our data, yet we make it anyway in an effort to see how this hypothetical would affect replacement decisions and expected profits. The dotted red line in the right hand panel shows our assumption about the required declines in daily rental rates. We have rental rates 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.*

We do not have the space to discuss the optimal replacement policy implied by scenario 2 in any detail, but 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 very adverse assumptions (which do not apply in our calculation of discounted profits under the *status quo,* the optimal replacement policy *still results in significantly higher profits than the status quo.* We conclude that our main find is quite robust to variations in our assumptions. Our estimates of the gains that can result from delaying the replacement of rental vehicles are likely to be extremely 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.

	Compact All Locations	Luxury All Locations	RV All Locations
Quantity	Value	Value	Value
Expected Discounted Values Under Optimal Replacement Policy			
$V(0,0,r_0)$	245680	337853	275614
$(1-\beta)V(0,0,r_0)$	20.19	27.77	22.65
$V(0,0,r_0)/\overline{P}$	25.4	14.4	14.7
Ratio of Expected Values: Optimal Policy versus *Status Quo*			
$V(0,0,r_0)/V_\mu(0,0,r_0)$	1.25	1.06	2.01

Table 6: Robustness Check of Profits/Returns: Optimal Policy versus *Status Quo*

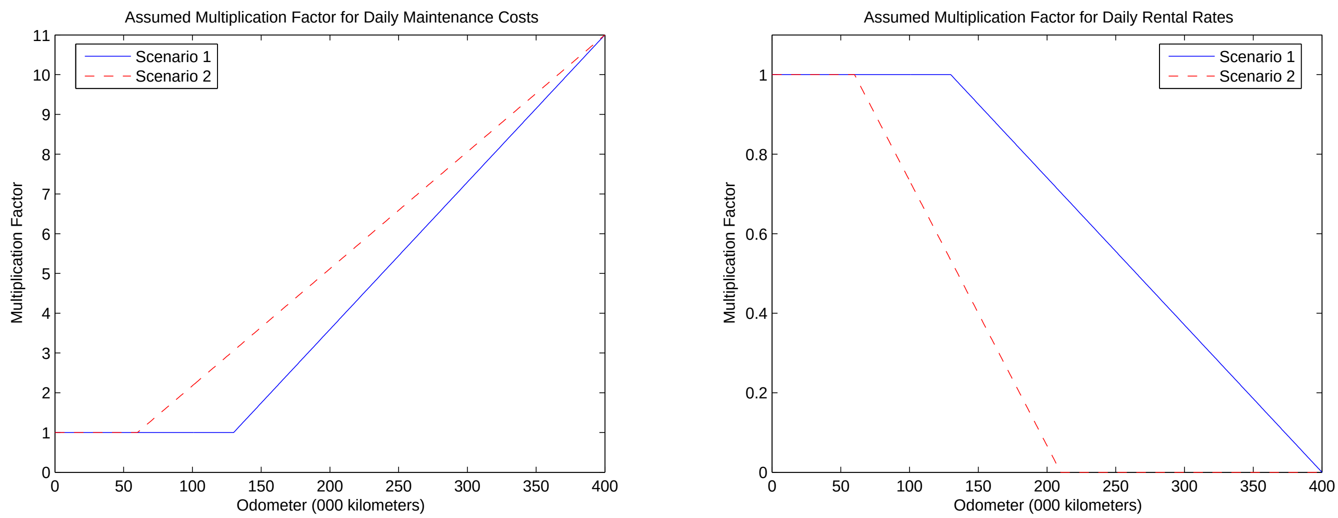

Figure 19 Multiplication Factors used for Robustness Check: All Car Types – All Locations

7 Extensions to the Analysis

Our analysis has been focused mainly on the question of the timing of replacement decisions, and we believe we have succeeded in providing convincing evidence that via modest changes in the company's operating strategy, it can significantly increase discounted profits. However our analysis leaves a number of unanswered questions:

1. Given how successful this company is at what it does, how could it fail to recognize the benefits from keeping its vehicles longer? Are there any overlooked constraints or regulations that might explain why the company decides to replace its rental vehicles "too frequently"?

2. Our analysis of sales prices revealed very large variations in the price received for apparently "observationally equivalent" vehicles. Why would the company "precommitt" to selling a vehicle on a particular date for the best price offered on that date, even if the best price seems below the fair market value for the vehicle?

3. Our analysis of the relative profitability of long and short term rental contracts revealed that for some vehicles, such as the compact car, short term contracts are significantly more profitable than long term contracts. Why should this be the case? If long term contracts are so much less profitable, why should the company allow its compact cars to be rented long term? Or alternatively, why doesn't the company increase the daily rates it charges for long term rental contracts?

4. Our analysis also revealed big differences in the overall profitability of different vehicles. In particular, the stream of discounted profits from rental of the RV or luxury car types are 20 and 40% higher, respectively. If these vehicles are so much more profitable, why not allocate more lot space on the margin to luxury and RVs, or alternatively, increase rental rates on compact cars to increase their relatively profitability?

These are some of the unanswered questions we have. As economists, we are accustomed to the standard sorts of "marginal arguments" for optimal decision making. Thus, this company must select a "portfolio" of vehicles for the lots in all of its 100 plus locations. Similar to standard portfolio analysis in finance, at an optimal allocation the company should be getting roughly the same expected "risk adjusted return" from an investment of $\$X$ in car type τ_1 as it does for an equivalent investment in car type τ_2, otherwise if there is one type of car that has a higher return per dollar invested, then the firm would be better off investing the marginal dollar in the car type that yields the highest possible returns.

Our analysis has revealed that of the three car types we have analyzed, the compact has the *highest rate of return* even though it has the lowest *discounted value of profits per car.* It is not completely obvious that the correct way to think about the firm's allocation problem as choosing to invest in the car with the high marginal return, or to allocate cars to a fixed level of lot space to maximize the overall value of discounted profits. These two, both quite compelling ways of viewing the overall allocation problem, seem to result in different allocations, at least on the margin. That is, if the company wants to get the highest return on its investment, it would appear it should allocate more of its vehicle "portfolio" to compacts and less to luxury or RVs. However if it is interested in maximizing the expected present value of profits, then it would appear that it should allocate more of its vehicle portfolio to the luxury and RV car types.

There could be complimentarities between cars of different types, and the firm should try to cater to its customers' preferences, and clearly some customers will want to rent compacts, others will prefer RVs and others will prefer to have luxury vehicles. If the company happens to be "stocked out" of a particular customer's most preferred type of vehicle, having a portfolio with sufficiently close substitutes may enable the company to keep that customer, as opposed to the customer walking down to the next rental company window to see if a competitor has their preferred vehicle in stock and ready to rent.

Our data does not include information on customers, their arrival rates to various rental locations and driving/return patterns (i.e. the probability that a car rented at location A will actually be returned to location B). Without more data on customer choices, and data on the company's competitors, it is difficult for us to formulate a more comprehensive model of the overall operations of this company. However we believe the analysis we have conducted in this paper constitutes a fundamental "building block" toward a more complete analysis of this optimal (i.e. profit maximizing) operation of this company. Whatever portfolio allocation of rental vehicles, and rental rates the company chooses, it will want to adopt a vehicle replacement policy that is optimal conditional on its vehicle portfolio and rental rate structure. For the moment, take the number of locations that the company rents vehicles from and the total space available for storing cars at each location as fixed (there is an even higher level decision that the firm faces about whether to close certain locations or open new locations, or expand at other existing locations that we ignore here).

Let M_i be the maximum number of cars that the firm has available in location i, $i = 1,\dots,N$. Suppose there are J possible car types (i.e. individual makes and models of cars), and the firm has adopted a *rental rate structure* $\mathcal{R}$ where initially we adopt the simplication that a rental rate plan for car type j at location i consists of two numbers $\{(R^l_{ij}, R^s_{ij})\}$ representing flat daily rental rates for long and short term rentals for each car type j at rental location i. Thus a rental rate structure consists of the complete array of all rental prices at all rental locations, $\mathcal{R} = \{(R^l_{ij}, R^s_{ij}), j = 1,\dots,J, i = 1,\dots,N\}$. Actually, the rental rate structure can become considerably more complicated than this once we consider generalizations of rental contracts that can include odometer-based discounts as we have considered in this analysis, and rental rates that depend on number of kilometers driven and not just a flat per day fee with unlimited kilometers.

Let $V_{ij}(\mathcal{R})$ denote the expected discounted value of profits from a car of type j in rental location i under the assumption tha the firm follows an optimal replacement strategy for each car type j at each location i under rental rate structure $\mathcal{R}$. Let $\overline{P}_j$ be the new purchase price of car type j. Then we can formulate the overall "optimal rental operations problem" as the following programming problem

$$\max_{\mathcal{R}} \max_{\{N_{ij}\}} \sum_{i=1}^{N} \sum_{j=1}^{J} N_{ij}[V_{ij}(\mathcal{R}) - \overline{P}_j] \quad \text{subject to: } \sum_{j=1}^{J} N_{ij} \leq M_i, \quad i = 1,\dots,N. \tag{27}$$

Nested within this problem is the regnerative optimal stopping problem, that we have solved in this paper, that delivers the value function $V_{ij}(\mathcal{R})$ for all car types at all of the firm's rental locations. This formulation of the problem assumes that the firm is not "liquidity constrained" and chooses a rental structure and vehicle portfolio allocation to maximize its overall net value of profits. With better data on all of the company's rental locations and customer data, it may be possible to solve this programming problem. But we note that there are still other complications. We have suppressed the dependence on the rental rate structures $\mathcal{R}_c$ of the company's competitors, $c \in C$. There is clearly a larger competitive game, and the firm's value and the optimal strategy for selecting its vehicle portfolio and rental rate structure will clearly depend on the portfolios and rental rate structures chosen by its competitors. Solving this overall competitive equilibrium problem remains a challenging area for future research.

8 Conclusion

We view this paper as providing both a practical and methodological contribution.

From a methodological point of view, we have shown how it is possible to integrate econometric duration models and (regenerative) optimal stopping theory in order to evaluate the profitability of the operating strategy of a firm. Further, we have shown how this apparatus can be used to test the hypothesis that the firm is a profit maximizer, and we have provided convincing evidence that the firm is not maximizing discounted profits.

The practical contribution of the paper is to provide both a framework and concrete computer code that enables us to characterize in a great deal of detail, the precise form of a profit maximizing replacement strategy for this firm. Our work can therefore result in actual application of theory that may have a concrete practical benefit to this firm and similar firms in the rental car business.

Our study relates to the recent econometric literature on "treatment effects" that attempts to predict the impact of various "treatments" that might be given to a "subject". In this terminology, the company is the "subject" under consideration, and the "treatment" we are suggesting is the strategy of keeping its rental vehicles longer before they are sold, combined with offering its customers appropriate discounts to induce them to rent the older vehicles in its fleet. The "treatment effect" we are interested in measuring is the increase in returns or discounted expected profits from adopting the suggested treatment.

In the treatment effects literature, the "gold standard" is to use controlled, randomized experiments to measure the treatment effect by comparing the outcomes for a randomly selected treated group to the outcomes for a control group. It would be possible for this company to conduct an experiment to evaluate the profitability of the alternative replacement and rental strategy we are suggesting, but it might difficult to design an idealized randomized controlled experiment where we can compare treatment and control groups at the same point in time. In particular, while it might be possible to randomly assign car ID numbers to those in a "treatment group" and those to a control group (i.e. the control group would be governed by the company's existing or *status quo* replacement and rental policy), there is a danger of various types of "contamination" between the treatment and control groups. In particular, if managers find out that the treatment effect is positive, (i.e. that replacing rental cars less frequently is indeed more profitable), they are likely to be tempted to also replace cars in the control group less frequently, and it this occurs, the measured "treatment effect" could appear to be zero. Similarly, rental customers who are randomly assigned to the treatment group (i.e. being offered the option to rent an older vehicle at a discount) may inform others that they know about this "deal" and other customers who are not randomly assigned to the treatment may request the treatment, and again this creates a temptation for managers to offer the same deal to customers in the control group.

For these reasons, different types of "quasi experiments" are likely to be undertaken. An example is to choose the treatment group to be all cars at a completely different rental location than the control group, or to focus on a given rental location and conduct a "before and after" test, by adopting the new replacement and rental strategy for a given make/model of rental cars at a specific rental location and compare the profits earned on the cars that were subject to the new replacement policy with the profits earned by the cars subject to the pre-existing *status quo* replacement policy. There are various sources of contamination in the measured treatment effects of both of these types of quasi experiments. In the before/after experiment, the treatment effects could be confounded by "macro/time effects," i.e. the economy and thus rental rates, rates of arrival of rental customers, maintenance costs and so forth could be different in the "after" period than in the "before" period, and so some of the measured treatment effects could actually be due to these other factors. Similarly, if the company were to conduct an an experiment in location B and compare the profits to those earned at location A (which continues to use the *status quo* replacement policy), the treatment effects could be confounded by "location effects" (i.e. different rental rates and arrival rates of rental customers, and different driving patterns and durations of rentals, etc.).

Some of these problems can be controlled for by adopting a third "blind control group" and using a "difference in difference" econometric strategy. Thus, for the example of a "before-after" experiment, while the treatment could be applied to a specific make/model of rental vehicles, the treatment could be withheld to another similar vehicle class which we could think of as the "blind control group". Then we could compare the difference before and after profit outcomes in the blind control group to the difference in before and after profit outcomes of the treatment group, and assume that the macro/time effects operate similarly for cars in both the treatment and blind control groups. Then any change in profitability in the blind control group could be ascribe to these macro/time effects, and thus, to the extent that the change in profits for cars in the treatment group was greater or less than this baseline, we would ascribe the difference to the "causal effect" of the treatment. Similarly, in the case where the treatment group is in location B and the control group is chosen in rental location A, we can choose a blind control group to be rental cars of a similar make and model that are rented at both locations A and B. Thus, differences in profits in the blind control group would capture location effects, and thus, the “treatment effect” would be measured by the extent to which the profits in location B exceed profits in location A, less the difference in profits of the cars in the blind control group in these two different locations.

This suggests that it is possible to “test” the predictions of our econometric model using experimental or “quasi experimental” methods. The use of blind control groups and difference in difference methods can help to control for some of the shortcomings of quasi experimental methods when true controlled randomized experiments are not feasible. However in either case, a major drawback of experiments is the time and expense of conducting them. The advantage of computerized models is that millions of “experiments” can be conducted under fully ideal conditions in a computer in a matter or minutes or hours, compared to many years that would be required to follow cars in a treatment group from new purchase, over their (longer) lifespans, until all of them were sold. However the drawback to computerized models is the issue of *credibility,* i.e. does the computerized model provide a sufficiently accurate model of reality to be trusted? We have attempted to answer this via the use of computerized simulations, by showing that our simulations provide close approximations to actual outcomes. However ultimately, the company (especially if it is risk averse) may want to conduct real experiments rather than rely completely on the predictions of a computerized model. However if the company finds these results sufficiently convincing to undertake experiments and other forms of data gathering necessary to evaluate its decision making and determine if there are profitable modifications it could make to its operating policies, then we think it could represent a “win-win” situation for science and for firms and other decision makers. To our knowledge, most companies do not have access to advanced scientific methods to evaluate and inform their decision making, and at the same time, economists rarely have access to firm data and the ability to suggest/conduct experiments that help them to improve and update their models.

Our hope is that this study will provide a practical example of how scientific methods can be of use, and the results will be sufficiently convincing to induce this company (and other companies) to make greater use of these methods and to undertake the additional experiments and data gathering necessary to validate the models and provide information about unknowns needed by the models to characterize optimal strategies and make accurate predictions of their actual effects.

References

[1] Cho, Sungjin (2006) "An Empirical Model of Mainframe Computer Investment" *Journal of Applied Econometrics* forthcoming.

[2] Howard, R. A. (1971) *Dynamic Probabilistic Systems: Volume II: Semi-Markov and Decision Processes* John Wiley & Sons, Inc. New York.

[3] Rothwell, G. and J. Rust (1997) "On the Optimal Lifetime of Nuclear Power Plants" *Journal of Business and Economic Statistics.* **15-2**, 195–208.

[4] Rust, John (1987) "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher" *Econometrica* **55–5**, 999-1033.

[5] 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.

[6] Rust, John (1996) "Numerical Dynamic Programming in Economics" in H. Amman, D. Kendrick and J. Rust (eds.) *Handbook of Computational Economics* Elsevier, North Holland.