Strategic timing of investment over the business cycle: Machine replacement in the US rental industry[†]

John McClelland
American Rental Association
John Rust
Georgetown University

October, 2016

Abstract

We analyze a large data set on rental revenues, maintenance costs, and sale prices to econometrically estimate key relationships needed to implement a dynamic programming model of the optimal timing of replacement of rental equipment owned by a large multi-location firm in the equipment rental industry. We find significant potential to improve rental company profitability from the strategic timing of equipment replacement, resulting in profit gains ranging from 1 percent to over 1700 per cent. The gains from the optimal replacement strategy come from exploiting seasonal variation in rental demand and the timing of the business cycle due to their effects on rental revenues and the cost of replacement. For some machines we find the optimal replacement strategy is pro-cyclical, but for others we find that a *countercyclical replacement strategy* — one where replacements are concentrated in slow periods of the business cycle — can significantly increase firm profits.

Keywords: rental equipment, rental markets, optimal replacement

[†]Direct questions to John McClelland at `John.McClelland@ararental.org`. The data used in this study are confidential and cannot be made available to third parties as per the conditions of a nondisclosure agreement between the data provider, the authors, and the American Rental Association.

1 Introduction

There has been extensive work in economics and operations research on optimal replacement of machinery. However the type of techniques that have been developed in academic settings have not been widely applied in the equipment rental industry until recently. This industry rents equipment used for earthmoving, high-reach access, material handling, portable electric power and many other construction and industrial uses.[1] The American Rental Association was founded to serve and represent the interests of firms in this industry in 1955, and has grown rapidly along with the industry itself. In 2015 the industry had over 5,000 firms operating in over 12,000 locations across North America with annual revenues over $45 billion and investment spending in excess of $12 billion, the majority of which is replacement investment by firms selling existing used machines and replacing them with new ones.

A small number of software and support firms have entered to provide sales and data support to the rental industry. The first entrants assisted rental companies in the sale of their used equipment, often via auctions. This support has expanded to warehousing the large, rapidly growing data generated from the hundreds of thousands of individual machines that are owned by the industry. This has lead to large data sets that can potentially be mined to help improve the operations of rental firms.

However the development of "rental analytics" – a combination statistical analysis of rental data with sophisticated computer modeling designed to improve the profitability of rental company operations — is relatively new. The analytical techniques used by the industry rely on calculations that define the current state of business operations without regard to how the business can be optimized over time. Similar to many other industries, equipment replacement and investment decisions are typically made informally based on the experience and intuition of seasoned managers and owners of equipment rental firms operating without the assistance any formal mathematical models. The goal of this paper is determine whether there can be value added from the application of the optimization techniques that are studied from a more abstract, theoretical level in academia, and whether sufficiently applied versions of these methods might be able to improve decision making and profitability in the equipment rental industry.

A realistic model of firms' optimal replacement decisions should account for various types of *uncertainty* such variability in rental revenues, maintenance costs, and the replacement costs. All of these are subject to idiosyncratic shocks that vary across machines and locations as well as "macro shocks" that

[1] For historical reasons, the definition of this industry excludes certain types of rental firms such as car rental companies, but includes other types of non-industrial rentals, such as tents and associated equipment for parties, events and weddings.

affect the overall demand for rental equipment. These shocks reflect business cycle fluctuations that can have especially high amplitude and persistence in the construction sector of the economy.

Replacement demand, and thus production of machinery to satisfy replacement demand, tends to be highly *pro-cyclical.* That is, replacement demand tends to be highest during "boom" periods when GDP growth is relatively high whereas replacement demand is lowest during "bust" periods when GDP growth is low or negative. Business cycles as well as higher frequency seasonal cycles affect rental revenues and investment in new machines by the equipment rental industry. Rental revenues and investment are particularly sensitive to fluctuations in the construction industry, especially non-residential construction spending. Demand for rental equipment is highly pro-cyclical, so replacement investment is pro-cyclical as well. This cyclicality can lead to high prices of machines and delivery delays for new machines during boom periods when many firms in the industry (as well as other construction firms that own rather than rent their equipment) try to replace their aging equipment at roughly the same time, often just at the start of boom periods.

In this paper we show that there may be profit opportunities from adopting a *counter-cyclical replacement policy* where more replacement investment is done during recessions or periods of weak demand for construction machinery. The weak demand for new equipment and relatively lower levels of replacement investment can reduce the price of new equipment while raising (at least in relative terms) the resale price of used equipment. This means it can be significantly cheaper for firms to replace machines in a recessionary period than in a boom period. In this paper we quantify the gains to following such a countercyclical replacement policy and derive conditions where it can be shown to be an *optimal replacement policy.*

Of course, the term "optimal" is relative to a number of important simplifying assumptions that make our analysis tractable. We abstract from a number of the larger, more difficult problems that rental companies face, including setting their rental rates, and the allocation of different machines in their (location-specific) rental equipment portfolios. In addition, we abstract from *capital constraints* that may be important drivers of pro-cyclical investment policies at many rental firms. Thus, our analysis needs to be qualified and interpreted relative to these simplifying assumptions. The policies we derive may not be optimal in the context of a more complex, realistic and encompassing model of a rental firm's operations that is able to account for all of these other features and constraints. However some of the larger publicly traded firms in the rental industry may not be liquidity constrained and may be able to follow a counter-cyclical equipment investment/replacement strategy. In addition, there may be certain aspects of tax policy that could make it profitable for some rental firms to follow a counter-cyclical replacement policy. Thus, our analysis might be considered as a sort of best-case analysis of the potential profit gains that might be achievable by firms that do not face strictly binding liquidity constraints and where the other simplifying assumptions in our analysis are not greatly at odds with reality.

Section 2 reviews some of the previous literature on dynamic optimization problems as they apply to optimal replacement of capital assets. Section 3 describes the data we were provided, a sample of equipment rental histories of specific machines from a fleet in six broadly classified geographical regions in the US. Section 4 provides an econometric analysis of these data, to enable us to predict rental revenues, maintenance costs, original equipment costs (OEC) and resale prices of machines, and also to characterize and predict the replacement policies used in the different rental locations in our sample. Section 5 discusses the specific dynamic optimization problem we have formulated and solved, and shows how the model can be used to analyze the profitability of different replacement strategies for equipment in the U.S. equipment rental industry. Section 6 presents our results and findings, showing the optimal replacement strategies implied by our econometric predictions of rental revenues, maintenance, and replacement costs, and compares these to the actual replacement strategies used by firms in our sample.

2 Existing Literature on Optimal Machine Replacement

Samuelson [1937] was among the first to study asset replacement problems, which he formulated mathematically as the problem of determining the optimal time to replace a piece of equipment that is subject to deterioration by maximizing the present value of net returns from an infinite chain of assets. That is, he assumed that whenever the current machine is sold, it is immediately replaced by a new one.

Samuelson's basic approach has been extended in various directions including allowing uncertainty, such as allowing for random failures in equipment or uncertain costs of repairing or maintaining equipment. The presence of uncertainty required more advanced optimization tools, including the use of stochastic dynamic programming which can solve a wide range of dynamic optimization problems involving sequential decision making in the presence of various types of uncertain "shocks" including machine failure, and macroeconomic shocks that affect overall demand for rental equipment.

Rust [1987] showed how the dynamic programming approach could be applied to real-world problems, under the assumption that *firms were behaving optimally*. Rust used this method to study the problem of bus engine replacement at the Madison Metropolitan Bus company, which periodically replaces bus engines with new or rebuilt engines to improve the reliability of its fleet of buses. McClelland et al. [1989] developed an optimal control model for assessing replacement policies when the assets can be rejuvenated. Both of these analyses are extremely relevant to equipment rental companies which are ultimately subject to the decisions that are made by key individuals (e.g. fleet managers) who, in addition to outright replacement, can consider alternatives such as overhauls, rebuilds, and reconditioning of machines that involve rejuvenation of varying degrees that could be more cost-effective than outright replacement.

Though the assumption that firms maximize profits is close to sacred in economics (along with the assumption that all consumers maximize utility), a growing line of work in *behavioral economics* has lead to an increasing recognition among economists that profit or utility maximizing behavior cannot be taken for granted. The dynamic optimization problems that real world firms and consumers are assumed to solve are highly complex and often very difficult, if not impossible to solve when sufficiently realistic versions are formulated mathematically. There is a problem, known as the *curse of dimensionality,* (a term coined by one of the developers of dynamic programming, Richard A. Bellman, is his early book on the subject, Bellman [1957]) that suggests that there are some dynamic optimization problems that cannot realistically be solved even using the most clever algorithms and fastest available computers. Even though Rust [1997] has shown that it is possible to break the curse of dimensionality in some situations using *randomized algorithms,* for dynamic optimization problems involving continuous decisions (e.g. how much to spend on a refurbishment) Chow and Tsitsiklis [1989] have shown the curse of dimensionality is present and cannot be "broken" regardless of the type of algorithm or computer that might be used to solve it. Thus, it seems quite reasonable that firms may resort to approximations or *rules of thumb* to circumvent the difficulty of solving complex dynamic optimization problems. The Nobel Prize winning economist Herbert Simon used the term *satisficing* to describe how firms and individuals may actually operate when confronting complex problems, an idea that originated in his seminal empirical study of firm behavior, *Administrative Behavior* (Simon [1947]).

However to the extent that firms do use rules of thumb and learn by "trial and error" how to improve their profitability over time, there may be an opportunity to use formal dynamic programming methods for sufficiently simple problems to help improve their profitability. Cho and Rust [2010] studied the operations of a large Korean car rental company using an approach similar to the one used in this study. Specifically, they compared the firm's actual replacement policy to the policy predicted from the solution to a dynamic programming problem. Cho and Rust [2010] found a puzzling result: *rental rates are flat,* i.e. they do not decline with age or odometer value. Car rental companies often justify such a policy by replacing all of their cars quite rapidly, often after only one year and with less than 20,000 miles. The Korean car rental company that Cho and Rust [2010] studied held their cars longer than most American car rental companies, selling them when they were on average 2.7 years old and 75,000 kilometers on their odometers. However Cho and Rust [2010] found the company could significantly increase its profits (more than doubling it for some makes/models) by keeping its rental cars roughly twice as long as the company kept them under its *status quo* replacement policy.

To address concerns that its customers would not rent cars that were "too old" Cho and Rust [2010] suggested discounting the rental rates of the older cars to induce customers to rent them. The company was initially skeptical that this alternative strategy could increase its profitability, but it decided to carry out a *controlled experiment* to test the predictions of the Cho and Rust [2010] analysis. The company chose four of its rural rental locations as "treatment locations" where the rental replacement and pricing policy suggested by Cho and Rust [2010] was put into effect. The operating profits in the treatment locations were compared with those in six "control locations" where the company's *status quo* replacement and rental pricing policy continued to be followed. The experiment revealed that the discounts for renting older cars were highly attractive to many of the company's customers, but did not greatly reduce rentals of new cars. This increased utilization rates and overall rental revenue from the older cars in the treatment locations by significantly more than it had expected. The unexpected increase in rental revenues combined with the substantial savings in replacement costs by keeping its rental cars longer demonstrated that the company could significantly increase its profits by keeping its rental cars longer combined with age-based discounts to incentivize its customers to rent the older vehicles in its fleet.

One primary question to be addressed in the initial stages of research is whether the replacement policies that companies use are approximately optimal. That is, are firms' decisions approximately the same as an optimal replacement strategy from a dynamic program whose goal is to maximize profits? If the decisions made by firms are based on rules of thumb and industry experience, there is a possibility that formal modeling and optimization of replacement decisions could identify profit opportunities for firms. However it is important to qualify that any computer model depends on a number of *assumptions* and if the assumptions are wrong, the predicted "optimal" replacement policy implied by these assumptions may not be optimal in practice. Thus, it is important to test the validity of these assumptions by comparing the profitability of firms under their *status quo* replacement policy to the counterfactual "optimal" replacement policy predicted by the computer model.

The first and least costly way to do this is via *stochastic simulations.* That is, we develop a computer environment that can simulate the evolution of rentals, replacements, and other variables for a hypothetical firm operating under its *status quo* replacement policy. These simulations can also allow for different simulated sequences of *macroeconomic shocks* that affect rental revenues, equipment resale prices, and replacement decisions by the firm. Then, *using the same sequence of macroeconomic shocks* we can simulate a *counterfactual replacement policy* such as the "optimal" replacement policy predicted by our dynamic programming model. For various scenarios and horizons, we can then determine whether the simulations of this optimal replacement policy really do result in higher simulated profits than the firm's *status quo* replacement policy.

Of course computerized simulations still depend on a number of assumptions and if these assumptions are wrong, even the simulated outcomes under the firm's *status quo* replacement policy might differ from the replacement decisions that the firm would actually undertake under similar conditions. Further, the environment in which the firm actually operates may be significantly more complex and reflect features that are not modeled or anticipated in the computerized simulations. Thus, the ideal way to test whether there are any real profit opportunities that are suggested from our analysis from adopting a counterfactual "optimal" replacement policy is to undertake a controlled experiment similar to the one described in Cho and Rust [2010]. If the controlled experiments reveal clear profit gains to adopting an alternative replacement policy (as they did in the Cho and Rust [2010] study), we can be much more confident that the profit gains would also be realized if the alternative replacement policy were implemented in practice.

Some of the profit gains we identify come from exploiting patterns in resale prices of machines over the business cycle, particularly the tendency for replacement costs to be lower during recessions. As we noted, the rental industry follows a pro-cyclical replacement strategy that tends to generate and reinforce these pricing patterns. However if enough firms were to adopt a countercyclical replacement strategy, this could impact new and resale prices of machines and in effect, "arbitrage away" some of the profit opportunities we identify in this analysis. Our analysis is therefore only fully valid in an environment where the number of firms that alter their replacement policies to take advantage of these "arbitrage opportunities" is not large enough to affect market prices. A more sophisticated analysis would be required to account for changes in market pricing patterns if we expect an industry-wide shift in replacement policies to occur.

3 Data

We were provided anonymized data on OEC (original equipment cost) and monthly rental revenues, and a separate sample of resale values (mostly from auction prices) and cumulative maintenance costs (from initial purchase until sale) for five types of rental equipment: 1) excavators, 2) high-reach forklifts, 3) scissor lifts, 4) skid steers, and 5) telescopic booms. The data are anonymous in the sense that though there are geographical identifiers for each of the machines in our sample, we were not informed of the identity of the company that owns this fleet.

For reasons we discuss below, we do not believe that our sample is representative of the overall rental industry in the US over the sample period January 1, 2011 to March 1, 2013. In particular, the data do not come from a random sample and the machines we study are not representative of the universe of all types of rental equipment. However machine types in our sample are the ones that rental companies own the largest numbers of, and thereby provide us a sufficiently large number of observations for econometric analysis.

Our data set has identifiers for six different broad regions of the US (identified as regions A to F) where the machines in our sample are located. In the rental revenue data set we have an equipment number identifier that enables us to track rental revenues earned by each individual machine in our sample, as well as its OEC and acquisition date and cumulative maintenance costs. We do have finer geographic identifiers: a variable that also indicates the specific urban area, but we do not know if a specific machines are "bound" to specific rental locations in these markets, or whether they are allowed to "float" between different locations in the same market area. Due to confidentiality concerns, we are unable to provide more specific information on the the makes of machines in our sample, or more specific information on the rental locations within the US where the machines are kept.

Table 1 summarizes the number of machines in each of the five categories in our equipment sample. We see that excavators are the most expensive of the five types of machines we analyze, and they are the youngest, with an average age of only 2.4 years. Not surprisingly, they earn the highest monthly revenue and have the second highest average monthly maintenance costs among the five types of machines in our sample. We have the most data on scissor lifts, and these machines are the cheapest in terms of OEC, and also earn the lowest average monthly revenue and have the lowest average monthly maintenance costs.

Table 2 summarizes the data in our sample of equipment sales. Since we have no information on whether any of the machines in our equipment sample had been sold during the period covered in our

Table 1: Summary of the Equipment Sample

Machine Type	Number of observations	Average age (years)	Average OEC	Average monthly rental revenue	Average monthly maintenance cost
Excavators	192	2.4	$141,042	$4,091	$144
High-Reach forklifts	581	4.1	94,325	2,645	150
Telescopic Booms	655	6.3	50,984	1,197	84
Skid Steers	307	2.8	22,806	992	65
Scissor Lifts	971	4.9	16,179	735	56

Table 2: Summary of the Sales Sample

Machine Type	Number of observations	Average age at sale (years)	Average OEC	Average monthly maintenance cost
Excavators	534	6.2	$116,423	$253
High-Reach forklifts	934	7.3	84,625	222
Telescopic Booms	996	8.0	47,992	109
Skid Steers	1607	6.2	18,378	108
Scissor Lifts	1494	7.0	15,539	60

equipment data set (or after it), we were provided a separate data set on resale prices (typically at auction) of a sample of 5565 other machines of the same five types as in our equipment sample that were sold between January 1, 2011 and November 27, 2013. We do not have any information on the rental histories or total rental revenue for these machines, however we do have the cumulative maintenance costs for each of these machines at the time they were sold.

Clearly the average age of machines that were sold are greater than an average age of machines in the current operating fleet. For example, the excavators in our sales sample were sold when they were on average 6.2 years old, whereas the average age of excavators in our equipment sample is only 2.4 years old. This difference reflects a wave of purchases of new excavators in 2012: only 47 of the 192 machines in our equipment sample were purchased prior to 2010, and 72 (or nearly 40% of the excavators) were purchased after January 1, 2012. This is in part a reflection of the phenomenon we noted in the introduction that rental companies tend to replace more of their rental equipment at the start of boom periods than during recessions.

In fact only 37 machines in our equipment sample were acquired during the depths of the recession

Table 3: Acquisition dates of machines in the equipment sample

	Number of machines acquired in calendar year										
Machine type	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012
Excavators	0	0	0	9	17	6	5	0	24	59	72
High-Reach forklifts	0	1	5	49	120	110	39	15	21	73	148
Scissor Lifts	0	0	11	56	232	330	110	17	24	63	128
Skid Steers	0	0	0	12	47	28	4	5	29	55	127
Telescopic Booms	13	32	38	94	179	174	63	0	9	8	41
Total	13	33	54	220	595	648	221	37	107	258	516

in 2009, whereas 516 machines were purchased in 2012, which nearly equalled the pre-recession peak in machine purchases in 2007 when 648 new machines were acquired. In general, there are relatively small numbers of acquisitions of all types of machines in 2002 and 2003 because machines acquired in those years would have been between 8 and 11 years old over the period of our equipment sample, and thus older than the mean age at which the company sells these machines. However in a steady state replacement cycle (i.e. in the absence of macro shocks and their effects on replacement decisions) we would expect to see between 12 and 17 percent of the machines in a company's inventory replaced each year. These are the oldest 12 to 17 percent of the machines the fleet, corresponding to mean times to replacement varying between 6 and 8 years.

For example if excavators are replaced every 6 years and we assume there are no cyclical effects and the demand for rentals is time-invariant, then in a steady state environment with no macro shocks affecting acquisitions or sales of machines we would expect that the average age of excavators would be about 3 years old and about 16% of these machines would be replaced each year (i.e the 1/6th of the inventory of excavators that would turn 7 years old each year). This implies that for the 192 excavators in our sample, we should expect to see about 24 machines replaced each year in steady state. If the rental companies in our sample followed a strict policy of replacing excavators when they reach 6 years old we would see no acquisitions prior to 2005 but observe acquisitions of 24 machines every year starting in 2005 and going forward, so that by 2011 (the first year of our equipment sample) the 24 machines would be 6 years old and due for replacement during that year.

While we do see very few replacements in 2002, 2003 and 2004, instead of a constant steady state flow of replacements we see a highly non-stationary pattern of replacement investment. The total number of replacements is well below the steady replacement rate level during the recession years prior to 2010, and is significantly higher than the steady state replacement level after 2010 after the Great Recession ended and the US economy was on a path to recovery. The same cyclical pattern of replacements is evident for all five the machine types in our equipment sample: relatively high levels of replacements (i.e. more than the expected number of replacements in a steady state situation) occurred in 2006 and 2007 (which were relatively strong years for the US economy), but there was pronounced collapse in the number of replacements in 2008 as the financial crisis struck the US economy. The low rate of new machine acquisitions continued into 2009 and 2010 during the depths of the ensuing recession. However as the US economy began to emerge from recession in 2011 and 2012 we see a rebound in the number of new machine acquisitions. The number of machines acquired in 2012 were nearly as large as the peak number of acquisitions prior to the recession. The main exception was telescopic booms: even by 2012 only 41 new telescopic booms were acquired, less than the steady state replacement level of 110 units, and far below the pre-recession peak replacement investment in telescopic booms of 179 and 174, respectively, in 2006 and 2007.

We used the monthly time series on US total construction spending (adjusted for inflation using the Producer Price Index) compiled by the US Census Bureau as our measure of overall construction activity that drives demand for rental equipment. Rental revenues are correlated with a construction spending index that we derived from this time series, and the correlation is most strong for excavators.

Though we are unsure of the precise method by which our sample was drawn, it does not appear representative of the growth in rental revenues for the rental industry as a whole over the period of our sample, and rental revenues and investment in new machines in our sample appear to co-move more strongly with overall construction spending than do rental revenues and investment for the industry as a whole. In particular we see significant seasonal and cyclical fluctuations in rental revenues and investment, and well as significant idiosyncratic variations in revenues that vary over firms, regions, and machines.

Even though our data may not be representative of the rental industry as a whole, we believe it has considerable value for analyzing the equipment replacement and investment decisions of the machines in our sample and assessing the extent to which there are potential profit gains from optimizing replacement decisions. Before we can do this optimization, however, we need to develop econometric models that can be used to predict and simulate rental revenues, maintenance costs, and the prices of new and used rental equipment — the topic of the next section.

4 Econometric analysis of equipment rental and sales data

This section summarizes the results of our econometric analysis of rental revenues, maintenance, and replacement costs. We estimated econometric models that capture the variability in the data that we noted in the previous section, and which accurately predict how these quantities vary by age and over the business cycle and during different months of the year, as well as across regions of the country. In addition, we used a separate sample on machine sales to predict the factors that affect when machines are sold. These econometric models are essential to our ability to mathematically model replacement policies and derive optimal policies by dynamic programming, as well as to simulate firm behavior under firms' *status quo* replacement policies.

4.1 Rental Revenues

We estimated regressions of rental revenues of the form

$$\log(R_{i,r,t}) = X_{i,r,t}\beta + \varepsilon_{i,r,t} \tag{1}$$

where $R_{i,r,t}$ is the rental revenue earned by machine i in region r at time t. If we assume that the unobserved error terms in this regression equation, $\varepsilon_{i,r,t}$, are normally distributed, this results in a "lognormal" model of rental revenues where predicted revenues are equal to

$$E\{R_{i,r,t}|X_{i,r,t}\} = \exp\{X_{i,r,t}\beta + \sigma^2/2\}, \tag{2}$$

where σ^2 is the variance of the residual unpredicted component $\varepsilon_{i,r,t}$ of monthly rental revenues.

The regression equation (1) was estimated for each of the five types of rental machinery in our data set (excavators, high reach forklifts, scissor lifts, skid steers and telescopic booms) separately, resulting in five separate sets of (β,σ^2) coefficient estimates for each of the machine types. We will be more specific about the $X_{i,r,t}$ variables entering the regression equation (1) shortly, but it includes fleet dummy variables, regional dummy variables, an index of construction spending, dummies for different makes of machines, monthly dummy variables, and the age of the machine, measured in months.

However there are some months where a particular machine will not be rented at all, so its revenues will be zero, $R_{i,r,t} = 0$. Since we cannot take the logarithm of 0, we estimated the regression equation (1) only for the subset of machines that earned some positive amount of rental revenue in a given month. As is well known, the predicted revenues from a regression that uses only positive values of the dependent variable will be subject to an upward "selection bias". That is, the regression will result in a prediction or *conditional expectation* that not only depends on $X_{i,r,t}$, but also conditions on the event that revenues are positive, which we can express as $E\{R_{i,r,t}|X_{i,r,t}, R_{i,r,t} > 0\}$.

We estimated a separate binary logit model to predict the probability that a particular machine would have zero rental revenue in a given month. Let $P(X_{i,r,t})$ be the probability that a machine i in region r with observable characteristics $X_{i,r,t}$ will have zero revenues in month t. We have

$$P(X_{i,r,t}) = \frac{\exp(X_{i,r,t}\gamma)}{1+\exp(X_{i,r,t}\gamma)}, \tag{3}$$

where γ is a vector of coefficients to be estimated (via the method of maximum likelihood). Then using the estimated β, σ^2 and γ parameters, we can write the conditional expectation (or best prediction) of rental revenues during the month as

$$E\{R_{i,r,t}|X_{i,r,t}\} = [1-P(X_{i,r,t}\gamma)]\exp\{X_{i,r,t}\beta + \sigma^2/2\}. \tag{4}$$

Intuitively, equation (4) adjusts for the upward bias in predicted rental revenues from the lognormal regression (1) that uses only data from machine-months with positive rental revenues by multiplying by the *between-month time utilization rate,* $1-P(X_{i,r,t}\gamma)$, which varies between 0.8 to 0.95, and is about 95 percent on average. Note that the other source of variation in time utilization, which we refer to as *within-month time utilization,* is captured indirectly via its effect on rental revenues in the regression (1).

In the interest of space, we omit the presentation of regression estimates of β and σ from the rental revenue regression equation (1), and instead we summarize the main conclusions we draw from these regressions and then present our predictions of revenues graphically for each machine type. One of the main findings is that expected revenues (conditional on the machine being rented during the month) decline with age of the machine.

The decline in revenues with machine age reflect two different effects: 1) a reduction in the rental rates earned by older machines, or 2) a reduction in the "within-month" time utilization of the machine with age. The latter effect refers to a decline in the days in the month an older machine is rented compared to a newer machine. Most of the expected revenue decline comes from the latter effect, since rental rates (whether measured at a daily level or on a per month basis) do not appear to decline with age. The decline in utilization with age may reflect a larger amount of time that older machines are off rent for maintenance each month, or a preference among customers (or rental company employees) for renting out newer machines instead of older ones when newer ones are available.

Figure 1: Unconditional expected monthly revenues, from equation (4), for machines in the sample

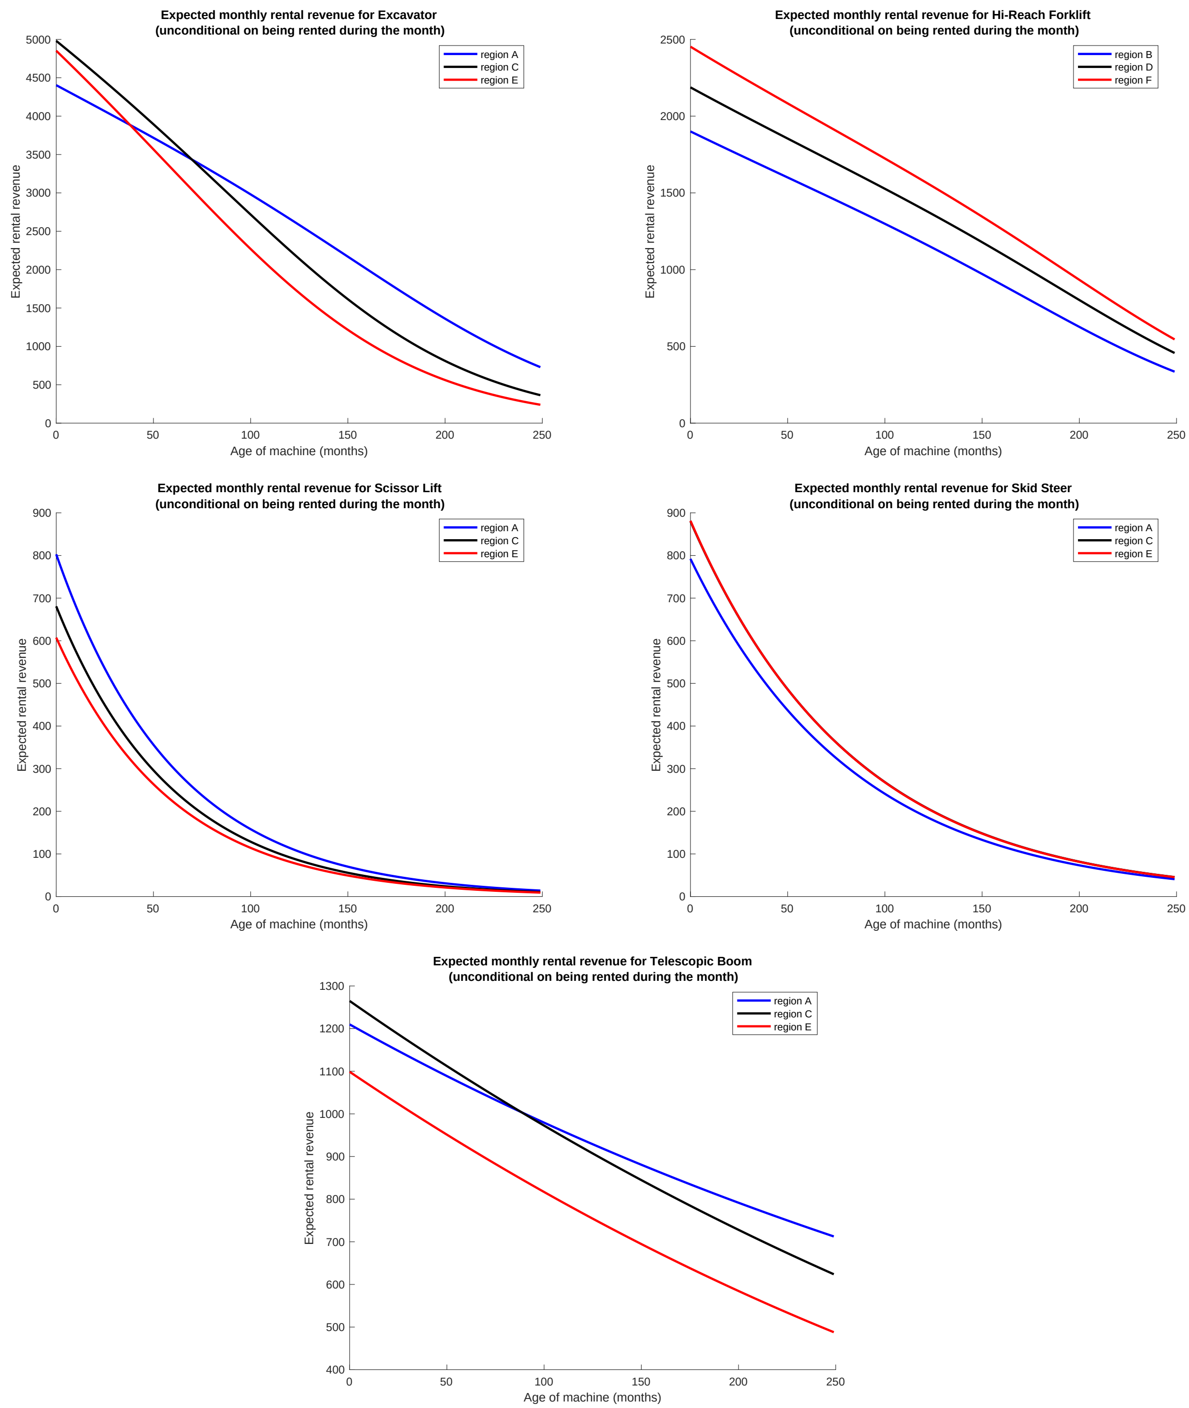

Figure 1 plots the expected revenue functions for a set of example machines and regions in our data set. In every case expected unconditional revenues predicted by our econometric model decline with the age of the machine. Revenues decline most rapidly with age for excavators and high reach forklifts, and this is driven mostly by the rapid decline in utilization rates as these machines age, because conditional their being rented during a given month, we find that the rental revenues of excavators and high reach forklifts do not decline rapidly with age.

4.2 Maintenance Costs

Our data also provided information on maintenance costs, but not on a month by month basis that would allow us to track variability over time in when maintenance costs are incurred, but only as a cumulative level of maintenance expenditures over the life of the machine. For the machines in our equipment sample we have total maintenance costs incurred from the date the machine was acquired until the end of our sample in March, 2013. In our separate data set on sales of machines, we also have the total maintenance costs from acquisition date until the date the machine was sold. That is, our data provide machine level data on the quantities $\overline{m}_{i,t,r}$ given by

$$\overline{m}_{i,t,r} = \frac{1}{t}\sum_{s=1}^{t} m_{i,s,r} \tag{5}$$

where $m_{i,s,r}$ is the actual maintenance costs incurred by machine i that is s months old in region r. Unfortunately we only observe the lifetime average maintenance costs $\overline{m}_{i,t,r}$ and not the individual monthly maintenance costs incurred in each month of a machine's life, $\{m_{i,s,r}|s = 1,\ldots,t\}$.

Our simulation and dynamic programming model will need to predict maintenance costs for each month in a machine's life, not just the average monthly maintenance cost over its entire life. However if we can predict how average monthly maintenance costs increase with age, we can also predict how expected maintenance costs increase with age in every specific month in a machine's lifetime. To see this, let the function $f_{i,r}(t)$ given by

$$f_{i,r}(t) = E\{\overline{m}_{i,t,r}\} \tag{6}$$

be the prediction from our regression model of average monthly maintenance costs of a machine that is t months old. Then $tf_{i,r}(t)$ is the predicted total or cumulative maintenance costs for this machine from acquisition to month t. Similarly, $(t-1)f_{i,r}(t-1)$ is the predicted value of cumulative maintenance costs from acquisition to month $t-1$, and thus the difference of these quantities,

$$E\{m_{i,t,r}\} = t f_{i,r}(t) - (t-1) f_{i,r}(t-1) \tag{7}$$

is the implied prediction of the expected maintenance costs $m_{i,t,r}$ in region r will incur on a machine i that is t months old *during month t*. It is easy to see that average monthly maintenance costs are an average of the presumably low maintenance costs when a machine is new and the higher maintenance costs when the machine is older, we will have $E\{m_{i,t,r}\} > f_{i,r}(t) = E\{\overline{m}_{i,t,r}\}$, and expected maintenance costs incurred in each month rise more quickly with the age of the machine, causing average monthly maintenance costs to increase with the machine age as well.

Figure 2 plots our estimates of current monthly expected maintenance costs, $E\{m_{i,t,r}\}$ and average monthly maintenance costs over the life to date of the machine, $E\{\overline{m}_{i,t,r}\} = f_{i,r}(t)$. We see that expected monthly maintenance costs (indicated by the blue curves in figure 2) do indeed increase significantly faster with the age of machines than the lifetime average maintenance cost (indicated by the red lines in the figure). Monthly maintenance costs are highest for excavators and high reach forklifts: a 100 month old excavator and high reach forklift are predicted to cost over \$1000 and \$750 per month, respectively, to maintain. In comparison, a 100 month old skid steer, scissor lift, or telescopic boom are predicted to cost less than \$250 per month to maintain.

With our estimated maintenance costs and expected revenue predictions, we can generate predicted expected profits for different machines, locations, months, and macro states. Let τ denote the type of machine (including the specific make, such as a make 1 excavator) and let r index the region of the country, m index the current month of the year, a index the age of the machine in months, and s index the macro state (boom, normal or bust). Our econometric model allows us to predict how revenues for these machines depend on each of these variables. Let $E\{R_i|a,\tau,r,m,s\}$ denote the expected revenues for a specific machine i that is a months old, of type and make τ in region r during month m in macro state s. Similarly, $E\{m_i|a,\tau\}$ denotes the expected maintenance costs for this machine. Note that our regressions for maintenance costs exclude month effects and macro shocks as well as region dummy variables, so the variables (r,m,s) do not enter the expected maintenance cost function. Let Π_i denote the gross profits earned by machine i in a given month, i.e. the difference between rental revenues less maintenance costs for a specific machine i with characteristics (a,τ,r,m,s). Then the expected profits for this machine is given by

$$E\{\Pi_i|a,\tau,r,m,s\} = E\{R_i|a,\tau,r,m,s\} - E\{m_i|a,\tau\}. \tag{8}$$

Figure 2: Current Monthly versus Average Monthly Maintenance Costs

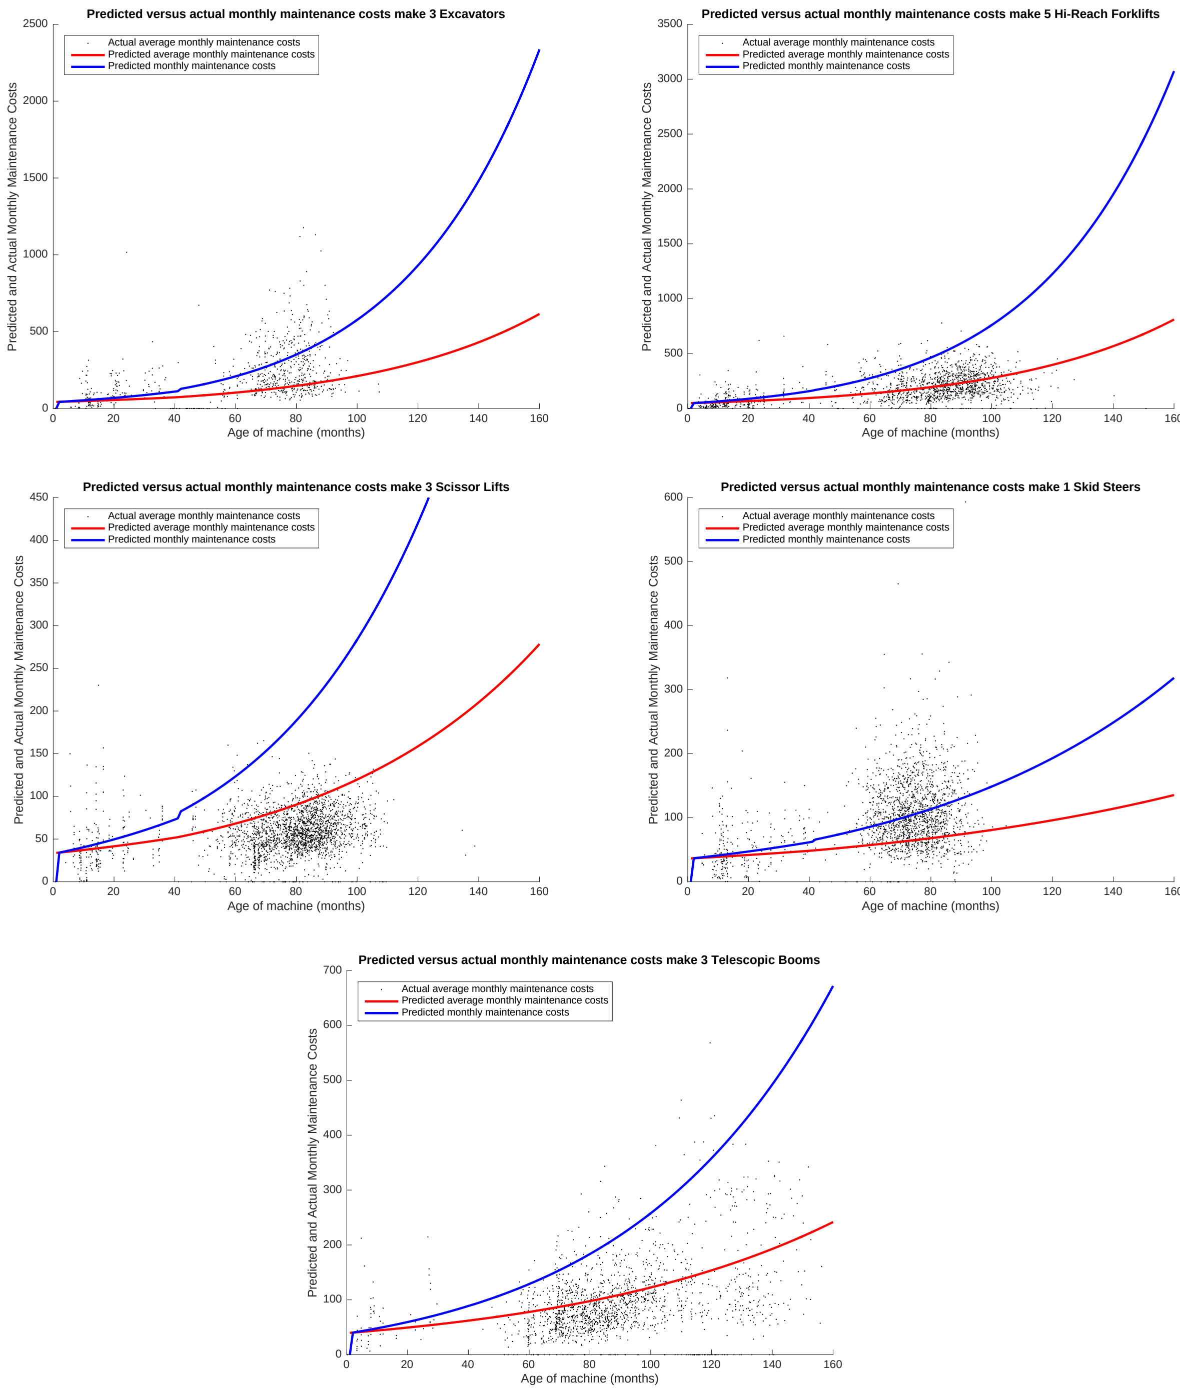

Figure 3: Expected Monthly Profits by Age of Machine

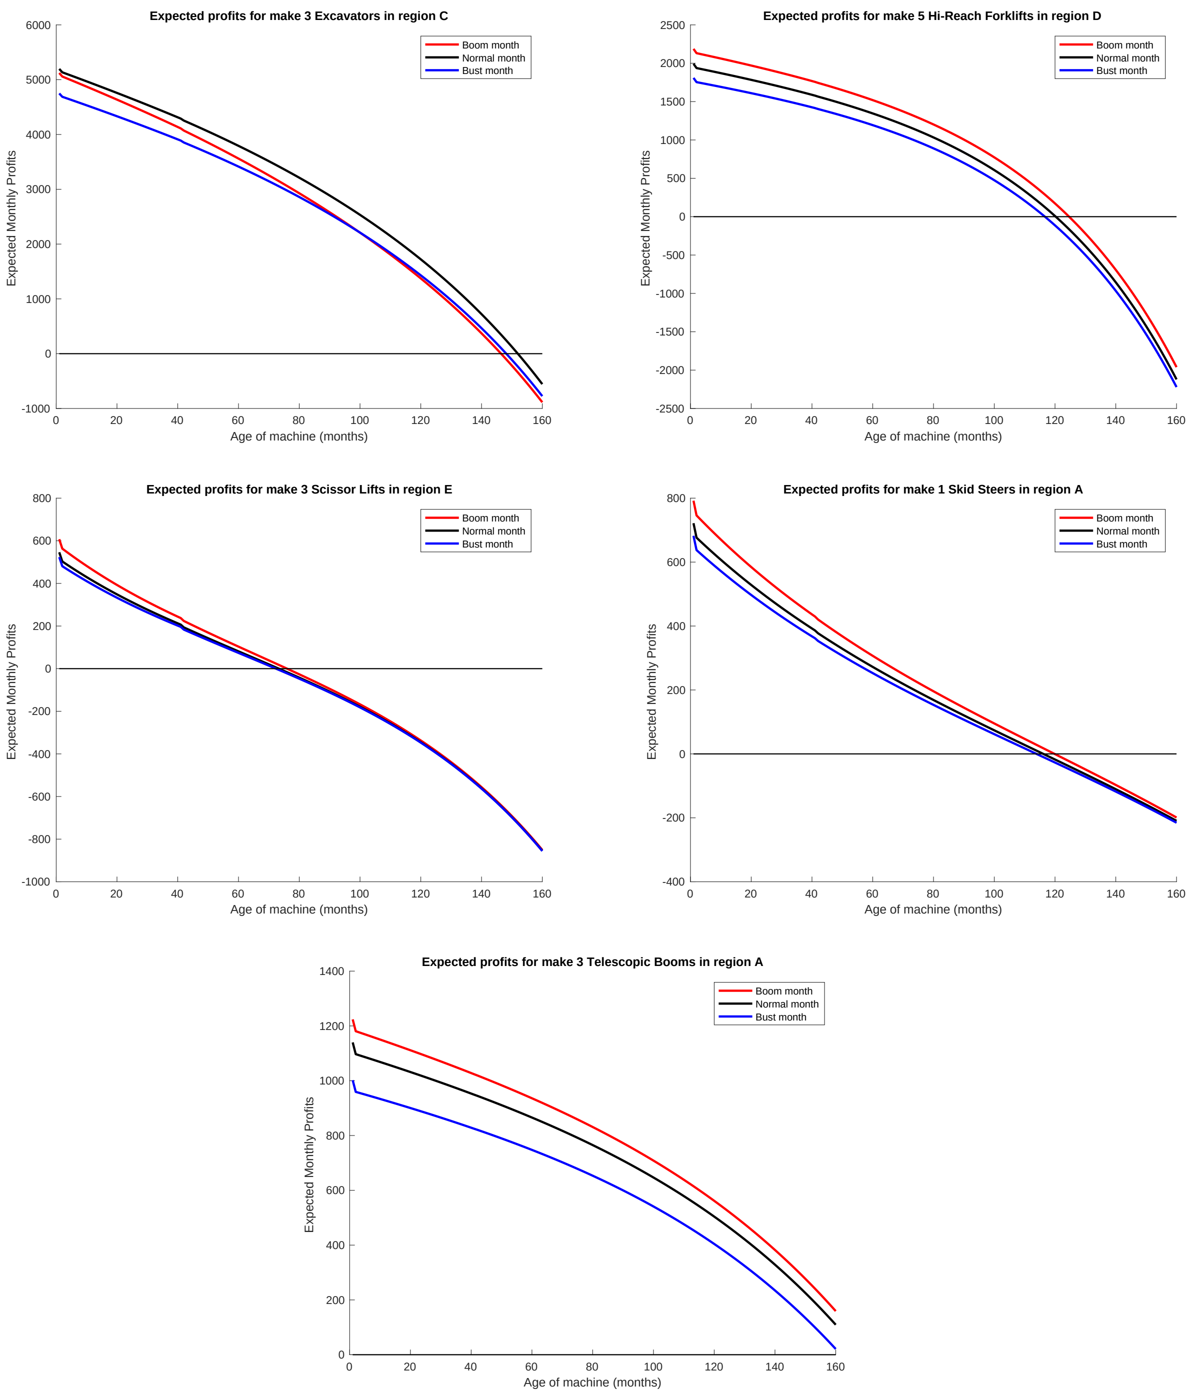

Figure 3 plots the expected monthly gross profits $E\{\Pi_i|a,\tau,r,m,s\}$ as a function of machine age a for different example machines (τ,r,m,s) where we fix the month as January ($m=1$). We see that profits are generally highest in the boom months ($s=3$, indicated by the red lines in figure 3) and lowest in the bust macro states ($s=1$, indicated by the blue lines in the figure). We also plot a horizontal zero profit line, so the age where the profit curves intersect the zero profit line indicates the "breakeven age" beyond which the firms can expect to lose money if they continue to hold the machine.

We see that the profits from renting a new excavator are the highest of the machine types shown. We forecast that a new make 3 excavator in region C will earn about \$5000 in gross rental profits in its first month of life if the macro state is in a boom ($m=3$) and about \$4500 if it is a bust state ($m=1$). The breakeven age for this machine ranges from 140 to 150 months depending on the macro state. That is, our model predicts that this firm will expect to lose money if it keeps this machine beyond this breakeven threshold. In the next subsection we econometrically model the equipment sales decisions of the machines in our sample. Our econometric model predicts that the company will have replaced this machine well before it hits this breakeven threshold: in a boom state, the median age at which we predict that the company will replace this machine is 75 months (i.e. there is a 50% probability that a 75 month old make 3 excavator will be replaced in a boom state), and the median age of replacement in a bust or normal macro state is about 100 months old. By the time the machine reaches the lowest breakeven threshold of 140 months, the probability this machine will be replaced is very close to 1.

The middle left panel of figure 3 shows that a make 3 scissor lift in region E is less profitable than the excavator that we we considered above: the monthly gross profit for a new make 3 scissor lift ranges from about \$600 in a boom state to \$500 in a bust state, and this profit decreases rapidly with age, and the breakeven point where gross expected profits turn negative occurs when the machine is about 70 months old. As we show in the next section, our econometric model predicts that the company keeps this scissors lift well *beyond* this breakeven age range: the median age of replacement in boom states is approximately 120 months and in bust states, 140 months. Thus, if our projections of expected revenues and maintenance costs is correct, our econometric analysis already suggests that the company could increase its profits by replacing this scissor lift substantially *earlier* than it currently does. In the next section we show that machines should be replaced *well before* they reach the current profit breakeven point. To determine these optimal thresholds an important additional consideration must be factored into the calculation: namely, how *replacement costs* vary as a function of the age when a machine is replaced.

4.3 Machine OEC and Resale Prices

The increase in maintenance costs with age combined with the decline in rental revenues with age are the two obvious factors that motivate rental companies to sell their older machines and replace them with new ones. But there is a key missing piece of information that helps firms determine the best time to do this: namely the expected cost of replacement. These costs depend on predictions of OEC for new machines the company buys and the secondhand resale or auction prices of the older machines the company sells. In this section we present our econometric predictions of these prices based on our OEC data.

We estimated a logarithmically transformed regression of the form

$$\log(P_i) = X_{i,t,r}\alpha + \varepsilon_i \tag{9}$$

where P_i is the OEC of machine i for the observations in our sample of purchases of new machines, and P_i is the resale (possibly at auction) price received for the machines that were sold. Separate regressions were run for the five different types of machines in our sample and the regression coefficient estimates of α and the estimated residual standard error $\sigma = \sqrt{\text{var}(\varepsilon_i)}$. We use $X_{i,t,r}$ to denote the explanatory variables in the regression, which include the age of the machine t, dummies for the region r, and make of machine, and dummies for the macro state s, the month of the year m. We initially estimated a simple specification that included a simple linear term in the age of the machine, a. This implies a constant rate of price depreciation for machines over their lifetimes. However we found we could significantly improve the fit of the model by allowing depreciation rates to vary with age. We estimated a spline specification that depreciation rates to differ depending on whether a machine is aged 0 to 40 months, and another depreciation to hold for machines over 40 months old. We also interacted these depreciation with dummy variables for the macro state being either a normal month ($s = 2$) or a boom month ($s = 3$).

The predicted OEC and resale values from our estimated regression model 9 are illustrated in figure 4. Focusing first on predicted OEC, we see that they are predicted to be higher in boom months in the case of telescopic booms, high reach forklifts, and scissor lifts, but *lower* in boom months in the case of excavators and skid steers. Turning to resale prices, the model predicts rapid early price depreciation in the case of excavators, high reach forklifts, and telescopic booms, but in the case of scissor lifts our model predicts depreciation rates that are *lower* for machines that are under 40 months old than for machines older than 40 months old. Also, our model predicts that the OEC of skid steers in boom months are *lower* than in normal and bust months, and this results in a lower price depreciation rate for skid steers under 40 months

Figure 4: Predicted OEC and resale prices for the 5 machine types, unrestricted specification

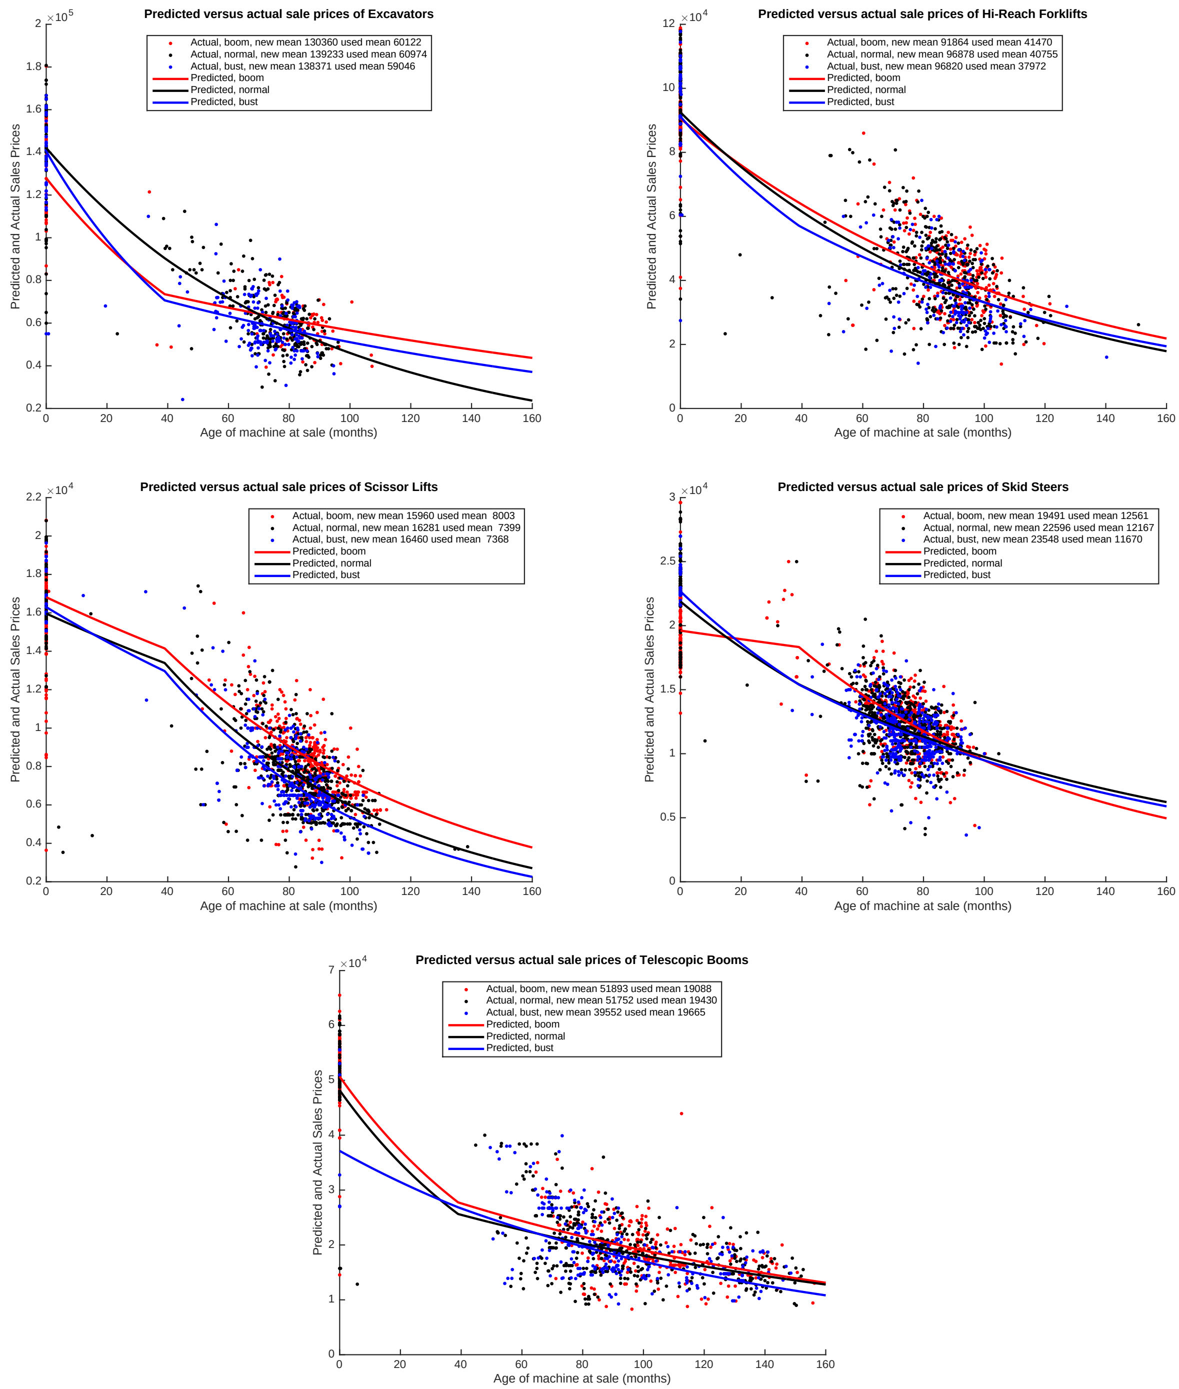

old relative to machines that are over 40 months old in boom months.

Our regressions had relatively small number of observations (ranging from a low of 760 observations for excavators to a high of 2320 observations for scissor lifts) and so some of our OEC and resale price predictions may be affected by outliers or machines that sold for especially high or low values due to special configurations or other idiosyncratic reasons. Thus, we advise caution in the interpretation of some of our predictions of optimal replacement strategies that result from these estimates, since the dynamic programming and simulation models treat these as providing reliable predictions of the prices it can buy and sell machines of different makes under different macroeconomic conditions during different months of the year. We would feel more comfortable if we had more data on OEC and resale prices to determine if our predictions are reliable.

4.4 *Status Quo* Machine Replacement Policy

We conclude this section by econometrically estimating a model of the decision of when machines are sold under the *status quo* replacement policies of regional fleet managers. There may be many idiosyncratic, unobserved factors that affect precisely when a fleet manager decides to sell a particular machine, but is seems clear that the age of and make of machine as well as the macro state s will be among the most important observable factors that affect these decisions. To this end we estimated a binary logit model of the decision of when to sell particular machines. Let $P_s(X_{i,a,r})$ be the probability that a machine i that is a months old in region r with observable characteristics $X_{i,a,r}$ will be sold. Under the logit specification we have

$$P_s(X_{i,a,r}) = \frac{\exp(X_{i,a,r}\delta)}{1+\exp(X_{i,a,r}\delta)}, \tag{10}$$

where δ is a vector of coefficients to be estimated (via the method of maximum likelihood).

It is important to understand that the data set used to estimate the δ coefficients is a *pooled dataset* that follows the rental histories of 2706 distinct machines provided in the "equipment sample" plus 5565 additional machines that are part of a separate sales sample. It would be impossible to estimate the δ coefficients using the equipment sample alone since the machine lifetimes (i.e. the duration from initial acquisition until sale) are *censored* — that is, all of the machines in the equipment sample were operating and none had been sold at the end of the 27 month time interval on March 1, 2013 in the equipment sample. However we were given a separate data set of 5565 machines that were sold between January 1, 2011 and November 30, 2013. This latter data set is uncensored in the sense we see the sales date for every one of the machines in this sample, but we have a potential problem of *choice based sampling* — i.e. the sales sample was specifically drawn to provide observations on machines that had been sold. The sales sample is also *left censored* — i.e. we do not observe the full rental histories for these machines as we do for the machines in our equipment sample.

However the sales sample does contain the *acquisition date* for each of the 5565 machines that were sold, and this enables us to *backcast* and reconstruct some information about these machines over their entire lifetime, except for rental revenues. For example knowing the acquisition date, we can determine information such as the month of the year and the value of Census construction spending index (which we have used to construct our trichotomous macro state variable s) in every month of the machine's lifetime between its acquisition date and sale date. Obviously in every month before the machine was sold, we know that it was not sold. Thus, our backcasted sales sample is more analogous to a separate panel data set on histories of machines, albeit one that has been *endogenously sampled.* However we believe the bias in our results by simply pooling the data sets and treating them as an *exogenously stratified sample* is small — we get reliable (i.e. unlikely to be badly biased) estimates of δ.[2] Due to space constraints, we omit the presentation of the estimated δ parameters but illustrate our estimates graphically when we compare the firms' *status quo* replacement policies to the optimal replacement policies that we calculate in secion 5.

5 Model

In this section we formulate a mathematical model of *optimal replacement of rental equipment.* We will show how to derive *optimal rental policies* — i.e. the dynamic strategies that specify when a company should replace individual rental machines in order to maximize the expected discounted profits from owning *an infinite sequence of machines.* That is, our mathematical model will tightly link *sales of an old machine with the purchase of a new replacement machine.* We assume that for each of the of the five types of machines in our sample, there is a stable, predictable "core demand" for these machines, so it is reasonable to assume that for these core machines, when there is a replacement of an existing old machine it will be replaced with a new one of similar make/model/configuration. Of course regional fleet managers can allow the inventory of machines of a given type to go up and down with business cycle conditions. In

[2] We will investigate the possibility of estimating δ using a modified likelihood function that accounts for the choice-based nature of the sales sample in future extensions of this work.

particular, they may hold more machines of a given type in boom times when demand is strong and hold fewer machines in bust times when demand is weak.

It is possible to extend the model we develop below to acommodate situations where a regional fleet manager may sell an old machine but not replace it (or delay when it buys a replacement), as well as situations where a rental location that is experiencing strong growth buys new machines without selling other older machines, resulting in a model with a variable total fleet size. However to accurately model such situations we would need a way to predict total demand for machines of each type at specific rental locations. If we can predict the demand for machines of each type (and possibly more detailed predictions by make/model and age of machine), it would be feasible to solve a more general version of the "replacement problem" that could enable us to determine a dynamically optimal policy for adjusting *fleet size.* Intuitively, in periods where we predict demand for machines will exceed the available stock of machines in "inventory" a fleet manager may be motivated to buy additional machines for a temporary period of time to meet the excess demand by its rental customers. Conversely in hard economic times, if the fleet manager forecasts a sustained period of low demand (and thus excess supply of machines of a given type), it may be optimal for it to "de-fleet" by selling some of its older idle machines.

We believe that important insights can be gained by considering initially a smaller "subproblem" of the overall fleet and portfolio optimization (and the related problem of setting rental rates for machines). If the overall management of a firm is approximately optimal (i.e. profit maximizing), it should not be possible to obtain significantly more profits by varying the management strategy of various subproblems of the firm's overall optimization problem. Conversely, suppose we find that there are ways to change policies the firm currently uses for specific well defined subproblems (such as when a firm should sell an existing machine and replace it with another one), that results in significantly greater profits. Provided this alternative strategy does not significantly impinge on or constrain or change other parts of the firm's operations and operating policy, it makes sense for the firm to consider adopting the improved policy since it results in greater expected profits even though the change may not be enough on its own to result in a "globally optimal" solution to the overall management strategy for a fleet at a given rental location.

Recall the three main *state variables* in our model (a,s,m), where a denotes the age of the machine in months, s denotes the "macro state" which takes the values 1, 2 or 3 where 1 denotes a recession or low demand month for the construction sector (which we have referred colloquially as a "bust month") when the firm expects rental demand for its machines to be low too, and 3 denotes a month when construction demand and therefore rental demand is at its peak, which we have described as a boom period. If a month is not a bust or boom month we refer to it as a "normal month" during which the firm can expect its "normal" demand for rental equipment.

The variable m denotes the current month of the year (Jan,Feb,...,Dec). We include this to capture seasonal effects on rental demand, such as the fact that demand for construction equipment is typically higher during the summer months than in the winter months, a pattern we typically found in our econometric analysis of rental revenues and rental equipment utilization in section 3. We allow the macro state variable s to evolve over time conditional on its past value and the month of the year. That is, we assume that $\{s_t\}$ evolves as a stochastic process, where the (discretized) value of the construction spending index $s_t \in \{1,2,3\}$ at each time t. We assume this process is Markovian with transition probability $\pi(s_{t+1}|s_t, m_t)$ where π is a transition probability that specifies whether the probability that the macro state variable s_{t+1} will take the value 1, 2 or 3 in period $t+1$ conditional on its value s_t in period t, and also on which month time t corresponds to, m_t. Obviously the month variable cycles deterministically according to the law of motion $m_{t+1} = \text{mod}(m_t + 1, 12)$ according to modulo arithmetic.[3]

It is obvious why we want to keep track of the age of the machine: our econometric analysis shows that in general rental revenues for a machine decline with its age whereas maintenance costs increase with age. Because of these two effects, it is generally the case that there will be a specific age or *replacement threshold* beyond which the firm will want to replace an existing old piece of rental equipment by a new (or newer one). We include the macro state s and the month of the year m as additional state variables for the obvious reason that rental revenues vary over the business cycle and over months in the year, since in the construction industry the peak period of demand is in the summer months, so construction demand for many types of rental machines, and thus rental revenues, tend to be higher in those months.

The OEC and cost of used equipment also varies with the business cycle and may also vary over different months in the year. In general, whenever we can econometrically uncover a solid relationship between our state variables and the relevant revenues, maintenance costs, and cost of new and used equipment, we want to include the key variables into our optimization model that help predict variations in revenues and costs that affect company profits. Our model can then adapt an optimal replacement policy to exploit or "arbitrage" these predictable variations and relationships in OEC, used machine costs, rental revenues, and

[3]Paarsch and Rust [2009] introduce a "cyclic inversion algorithm" to substantially speed up the solution of dynamic programs that involve stochastically cycling state variables such as s_t and deterministically cycling state variables such as m_t. We can use this algorithm to speed up the numerical solution of the dynamic programs we describe below.

maintenance costs to help firms earn greater profits.

There are other "implicit state variables" in our model that are geography based: we have data on rental revenues, maintenance costs and the costs of used equipment in different regions in the country and there can be different statistical relationships governing rental revenues and resale prices that hold in widely varying parts of the country, say between the Southwest and the Northeast. These differences can persist due to the transportation costs of moving heavy equipment across the country, and thus there can be differences in rental revenues and even in the costs of used rental equipment that will not be arbitraged away by moving rental equipment from low profit regions in the country towards high profit regions.

Let $V(a,s,m)$ be the expected present discounted value of profits of a given make/model of rental equipment at a given rental location in a specific region of the country. In terms of our notation in section 3, we have suppressed τ which indexes the type and make of a specific machine to simplify the equations of the model. We have also suppressed the region index r, but all of our results provided in the next section will account for different machine types, makes and regions. That is, in our analysis the value function and decision rule implied by the model also depends on (τ,r) in addition to (a,s,m), however the variables (τ,r) are assumed to be *time invariant* whereas the variables (a,s,m) change over time, something we denote by time subscripts (a_t,s_t,m_t) which denotes a machine of age a_t, in macro state s_t in month m_t at time t.

There is a functional equation for V known as *Bellman's equation* that we solve in order to determine the optimal rental equipment replacement policy for the firm

$$\begin{aligned} V(a,s,m) \quad = \quad \max \Bigg[& R(a,s,m) - C(a) + \beta \sum_{s'=1}^{3} V(a+1,s',m'), \\ & \max_{a' \geq 0} R(a',s,m) - C(a') + P(a',s,m) - P(a,s,m) + T(a) + \beta \sum_{s'=1}^{3} V(a',s',m') \Bigg]. \end{aligned} \tag{11}$$

where m is the index of the month of the year and evolves according to $m' = m + 1 mod(12)$ as discussed previously.

Equation (12) states that $V(a,s,m)$ is the maximum of two different options: 1) keeping the existing machine of age a, or 2) selling the existing machine a and purchasing a new machine a', but choosing the *best such replacement age* a'. This is indicated by the second max operation in the right hand term of equation (12). The function $R(a,s,m)$ denotes the *expected rental revenues* earned from rentals of a machine of age a in macro state s in month m, and $C(a)$ denotes the expected costs of maintaining that machine. The function $P(a,s,m)$ is the *market price function* that specifies the price the fleet manager expects to receive by selling its existing machine of age a when the macro state is s and the month is m. When it replaces the old machine, it buys another replacement machine of age a' for a price $P(a',s,m)$ and of course this is a cash outflow for the fleet and is why this enters with a negative sign in equation (12). The symbol T denotes the expected value of any *transactions costs* associated with trading the current used machine of age a for a replacement unit of age a'. These costs could include transportation costs of hauling out the machine to be sold and bringing in the new replacement unit, any licensing or transfer costs, and any auction or sales fee to an intermediary that assists the fleet manager in making this transaction. Finally the symbol $\beta \in (0,1)$ denotes the fleet manager's *discount factor* which is related to the interest rate (or risk-adjusted "opportunity cost of capital") r via the equation

$$\beta = \frac{1}{1+r} \tag{12}$$

The discount factor can also incorporate a risk of bankruptcy or closure of a specific rental location provided we assume there is a constant probability p that this will occur. Then a component of the discount factor β will include a "survival probability" $(1-p)$ which corresponds to the probability that the rental location avoids "dying" in any given month, which would result in a permanent termination of further cash flows.

As we noted above, we solved the dynamic programming problem by solving the Bellman equation (12) but under the constraint that when the firm replaces an existing machine, it can only purchase a new machine. This amounts to the restriction that $a' = 0$ in equation (12), in which case $P(a',s,m) = P(0,s,m)$ is the OEC, i.e. the cost to the firm of purchasing a brand new machine to replace the existing used machine. When we relax this constraint, the resulting value function will only increase: this reflects that providing the firm with additional options can only increase and will never decrease its present discounted profits. But as we noted above, we decided to impose the constraint that a used machine must be replaced by a brand new one. We did this because we are not sure of how "thick" markets for used machines are, and because of our desire to avoid conclusions that might be regarded as too radical or unconventional to firms in this industry, which to our knowledge, do generally follow a strategy of replacing used machines with brand new ones.

The Bellman equation can be used to solve for the function $V(a,s,m)$ and from this function we can derive the *optimal replacement policy* $d(a,s,m)$ where $d(a,s,m) = a$ if it is optimal to keep the current machine, and $d(a,s,m) = a'$ if it is optimal to trade the current machine of age a and replace it with another machine of age a'. The equation for $d(a,s,m)$ is given by

$$d(a,s,m) = \begin{cases} a & \text{if } v(a,s,m,0) > v(a,s,m,1) \\ a' & \text{otherwise} \end{cases} \tag{13}$$

where a' is the age of the optimal replacement machine from the bottom formula in the Bellman equation 12 and $v(a,s,m,0)$ is the expected payoff to keeping the current machine given in the top expression of equation 12

$$v(a,s,m,0) = R(a,s,m) - C(a) + \beta \sum_{s'=1}^{3} V(a+1,s',m') \tag{14}$$

and $v(a,s,m,1)$ is the payoff to replacing the current machine with the best replacement machine a' given in the lower formula of equation 12

$$v(a,s,m,1) = \max_{a' \geq 0} \left[R(a',s,m) - C(a') + P(a',s,m) - P(a,s,m) + T(a) + \beta \sum_{s'=1}^{3} V(a',s',m') \right]. \tag{15}$$

Typically it will be the case that $a' = 0$, i.e. the firm replaces an old machine with a brand new one. In most of our results, we place an arbitrary restriction that when a firm replaces an old machine it *must* buy a brand new replacement machine. But this is a restriction we can relax. If there is a sufficiently "thick market" the fleet manager might be able to do better by buying a slightly used machine rather than a brand new machine, especially if there is rapid early price depreciation for rental machinery similar to what we observe in the market for used cars. However the key is whether there is a sufficiently active market in very new but not brand new machines. If not, it may not be feasible to buy slightly used machines that are only a few months old. This is part of the reason why we opted to constrain the choice of replacement to only $a' = 0$, i.e. to constrain firms to buy brand new replacement machines at OEC.

5.1 Valuing existing, potentially suboptimal replacement strategies

Let $\mu(a,s,m)$ be the probability a firm replaces machines under its existing or *status quo* replacement policy. We have already econometrically estimated these probabilities in section 4.4 above. We assume that the *status quo* replacement policies always involve replacing any machine that a fleet manager sells with a brand new model, i.e. $a' = 0$. We are interested not to know only the *behavior* corresponding to firms' existing replacement policies (i.e. which ages, months and macro states make the firm most likely to replace an older machine), but we also want to know the value of *expected discounted profits* implied by these policies. Let $V_\mu(a,s,m)$ be the present discounted value of gross profits (i.e. pre-tax and not accounting for any allocation of corporate fixed costs, etc. as discussed above). Then we have the following recursive linear equation for $V_\mu(a,s,m)$ which is analogous to the Bellman equation (12):

$$
\begin{aligned}
V_\mu(a,s,m) \quad = \quad & [1-\mu(a,s,m)]\left[R(a,s,m)-C(a)+\beta\sum_{s'=1}^{3}V_\mu(a+1,s',m')\right] \qquad (16)\\
& +\mu(a,s,m)\left[R(0,s,m)-C(0)+P(0,s,m)-P(a,s,m)+T(a)+\beta\sum_{s'=1}^{3}V_\mu(1,s',m')\right].
\end{aligned}
$$

By construction, we have $V(a,s,m) \geq V_\mu(a,s,m)$, i.e. the optimal value function $V(a,s,m)$ from the solution of the Bellman equation 12) is always at least as great as the value of any other replacement policy $\mu(a,s,m)$, given by the value $V_\mu(a,s,m)$. Thus, the difference $V(a,s,m)-V_\mu(a,s,m) \geq 0$ represents the gain in discounted profits to the firm from switching from its existing *status quo* replacement policy $\mu(a,s,m)$ to the optimal replacement policy $d(a,s,m)$. However it is important to keep in mind that $d(a,s,m)$ is only "optimal" in a restricted sense: we are taking as given the firm's policies over how it prices its rental equipment and its overall fleet management and rental location "portfolio allocation" — i.e. how it allocates the limited space and capital in a given rental location over different types, makes and models or rental equipment.

6 Results

In this section we solve the dynamic programming problem described in the previous section to determine optimal replacement policies, using the econometric predictions of rental revenues, maintenance costs, OEC and resale values that were presented in section 4. We describe the optimal replacement policies in some detail for machine types: excavators and telescopic booms. We compare our calculated optimal replacement policies to the *status quo* policies followed by this firm using stochastic simulations of our econometrically estimated replacement rules described in section 4.4. We emphasize that there is no single "optimal replacement policy" but rather we calculated optimal replacement policies that are individually optimized for each machine make and rental location. In the notation of the previous section, our calculated optimal replacement rules do not only depend on (a,s,m) (i.e. the machine age a, macro state s, and month m), but also on (τ,r) where τ indexes the machine type and make and r indexes the rental location.[4]

[4]Note that due to space constraints we do not provide results for the other machine types: scissor lifts, high reach forklifts, and skid steers. However we do find significant potential increases in profit to adopting an optimal replacement strategy for these machines as well, and can provide a detailed analysis for these other machines to interested readers on request.

6.1 Excavators

The first results we present are for excavators. Figure 5 displays the optimal replacement policy and the corresponding optimal value function for a make 4 excavator in region A. We also compare the optimal replacement policy to the one the firm actually uses, as well as the implied discounted value of profits that the firm can expect to earn from an infinite sequence of these machines under the *status quo*. The top panel of figure 5 shows the optimal replacement thresholds by month, for the three different values of the macro state variable. We see that there is some variability in these thresholds across different months of the year, and this variation reflects predictable differences in rental revenues and replacement costs in different months. Generally, the DP model predicts that the replacement thresholds are the lowest in boom months ($s = 3$) and highest in normal months. Thus, for excavators we find the optimal replacement policy is *pro-cyclical* but with one key difference. Interestingly, the replacement threshold in a bust month lies in between the thresholds applicable in a boom month and a normal month. That is, under the optimal policy the firm is least likely to replace a machine in a *normal* month rather than in a bust month.

Thus, if January was a boom month, the DP model predicts that it is optimal to replace the excavator once it is 60 months old or older. However in a bust month, the model predicts that the firm should not sell the machine until it is more than 150 months old, and in a normal month, the firm should not replace the machine until it is more than 190 months old. The middle panel compares the calculated optimal replacement policy with the policy the firm actually follows. We see that in all three macro states, the firm replaces machines when they are roughly 60 months old, with a median age of replacement ranging from about 57 months old in boom times to about 70 months old in normal or bust months. Thus, the firm is not varying its replacement policy over the business cycle as much as the DP model predicts is optimal in order to take advantage of predictable changes in replacement costs and rental revenues over the business cycle.

The third panel of figure 5 plots the value functions as a function of the machine age and macro state and compares them to the corresponding value functions for the company under its *status quo* replacement policy. We see that the value functions corresponding to the optimal policy are uniformly higher. Specifically if we compare the solid red line (the value of the machine under the optimal policy as a function of age when the economy is in a boom month) with the dashed red line (the value of the machine under the *status quo*), we see that the solid red line is uniformly higher at all ages. If we consider a brand new machine, i.e. at age 0, the value function is $v(0,3,1) = 435984$, i.e. the firm can expect to earn a discounted

Figure 5: Optimal replacement policy for a make 4 excavator in region A

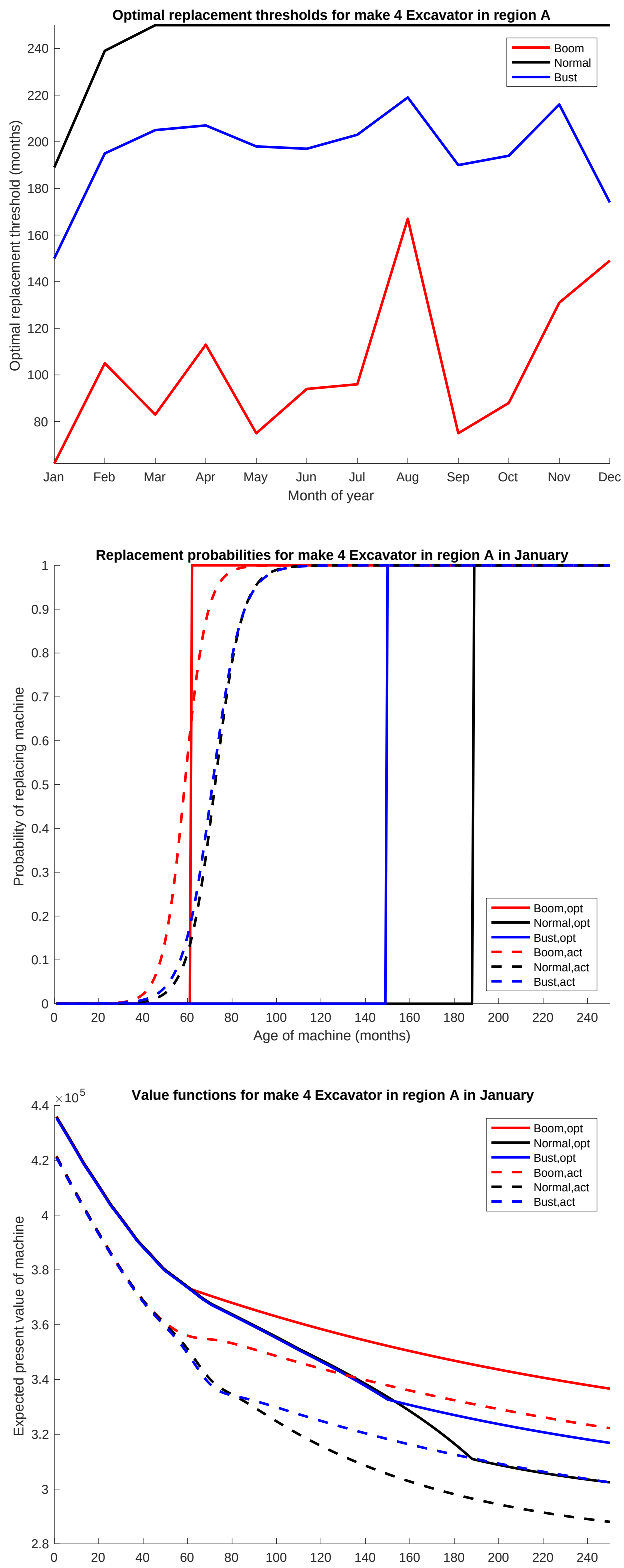

profit of \$435,984 *over an infinite horizon* by following the optimal replacement policy. This works out to a monthly gross profit of \$3618. In comparison, the expected discounted value of profits under the company's *status quo* replacement policy is $v_{\mu}(0,3,1) = 421528$, i.e. it expects to earn \$421,528 over an infinite horizon involving an infinite sequence of purchases and replacements of excavators. This works out to a monthly profit of \$3498. Thus, the optimal policy results in an increase in discounted profits of 3.4% ($0.034 = (3618 - 3498)/3498$).

To get some perspective on what we mean by "value", note that the average expected discounted profits the rental company expect to earn if it owns a new make 4 excavator (averaged over the different months and macro states) is \$436,519. It is important to realize that this number does not only include the discounted profits from the initial make 4 excavator from its initial purchase until replacement, but it also includes the entire stream of discounted profits from the sequence of all future make 4 excavators the company will purchase and rent to customers. We can calculate the expected discounted value of profits for the *current machine only* and is roughly half as large: \$232,449. The average of the predicted OEC costs to buy a new make 4 (again averaged over the different months and macro states) is nearly \$140,000. The valuations we calculated above presume that the firm already owned a brand new make 4 excavator. If it did not already own it, we would have to subtract the OEC cost required to buy a new make 4 in order to calculate the net profit to the fleet. Thus, the company expects to earn a gross discounted profit of \$232,449-140,000=\$92,448 from its ability to rent *this* make 4 excavator to its customers. However if we consider the company as infinitely lived with the ability to rent a sequence of make 4 excavators to customers extending into the infinite future, the expected discounted value of these profits are about 3 times the OEC of the initial make 4 excavator. This suggests that this particular location in region A is a relatively profitable one for the company because in effect it can expect to get more than a 3 to 1 return on its investment of about \$140,000 for its first of a sequence of make 4 excavators.

Of course, the use of an infinite horizon model involves some assumptions that one might question, such as whether the prices and rental revenues of make 4 excavators will remain the same over the infinite future. We can reasonably expect rapid technological progress and perhaps not too far in the distant future there may be new models of make 4 excavators that are fully *robotic* (i.e. they do not require human operators or can be directed and run remotely), and may have a number of other technological improvements that change their OEC and resale prices as well as the rental revenues the company can expect to earn. We have not attempted to adjust our predictions for such possible technological changes

Table 4: Detail on simulations of make 4 excavators in region A

60 month simulation	Discounted values			Undiscounted values		
20 replications	Optimal	*Status Quo*	Percent change	Optimal	*Status Quo*	Percent change
Revenues	\$231,439	\$238,260	-2.8%	\$293,320	\$301,712	-2.8%
Maintenance costs	14,246	8,951	59.3%	17,749	11,264	57.5%
Replacement costs	34,134	64,405	-47.0%	41,317	80,837	-48.9%
Gross profits	\$183,059	\$164,903	11.0%	\$234,254	\$209,611	11.8%

as they are too hard to predict. Our defense for failing to do this is that *since future profits are discounted, what happens far ahead in the future does not have a big impact on our calculations, which are dominated by profits received in the near term future which we can more confidently predict.*

In order to obtain more insight into how the optimal replacement policy enables the firm to earn greater profits and whether these increased profits can actually be realized in the near term we resort to the technique of *monte carlo* simulations. That is, we can simulate detailed future paths of revenues, maintenance costs, machine sales and purchases, as well as macro shocks that affect these quantities, under the optimal replacement policy and under fleet managers' *status quo* replacement policies. Further, we can do these simulations over shorter horizons to determine if the optimal replacement policy can actually deliver profit gains in a reasonably short period of time (as opposed to delivering most profits far off in the future when it is far less certain that the predictions of our model will still be valid). Further, the simulations provide more insight into actually how the optimal replacement policy can deliver improvements in profitability.

Table 4 provides a breakdown of average simulated gross profits for the make 4 excavators in one of its locations in region A into its key components: revenues, maintenance costs, and replacement costs. The table shows discounted and undiscounted values separately, and all are reported on a per machine basis, which is an average over all 50 simulated machines, each simulated for 60 periods, and finally averaged over 20 independent replications with different sequences of stochastic shocks such as macro shocks drawn in each replication. We see that because the optimal policy entails keeping these excavators longer than the company keeps them under its *status quo* replacement policy, revenues are lower and maintenance costs are higher under the optimal replacement policy than what the company would expect under the *status quo.* However *replacement costs* are the biggest expense for the company and by keeping these excavators longer and strategically timing their replacement over the business cycle, the optimal policy

Table 5: Gain in discounted profits from optimal replacement policy: excavators

Make/Region	A	B	C	D	E	F
Make 1	9.6%	2.7%	1.9%	5.1%	2.8%	3.2%
Make 2	4.7%	1.2%	1.2%	2.1%	1.8%	1.4%
Make 3	3.1%	1.4%	1.4%	1.4%	2.5%	1.4%
Make 4	3.5%	1.8%	1.6%	1.4%	3.1%	1.8%
Make 5	10.4%	2.9%	2.1%	5.5%	3.0%	3.5%

is able to reduce replacement costs by nearly 50%. As a result, the optimal replacement policy leads to an 11% increase in discounted profits and a 11.8% increase in undiscounted profits on average in our 20 simulation replications.

Table 5 reports the percentage increase in discounted profits from switching from the *status quo* replacement policies to the optimal policies for 30 different combinations of region and make of excavators. We solved separate dynamic programs that are specific to each region and make of excavator, and then simulated the optimal replacement policy over a 60 month (5 year) horizon for a fleet of 50 excavators. We initialized each simulation by drawing a set of initial ages of excavators that roughly match the age distribution of excavators at the beginning of our sample. Since the outcome of the simulations depend on the particular sequence of macro shocks that are simulated over the ensuing 60 month simulation horizon, the results in the tables are *averaged and reported on a per machine basis* and the averages are taken over *20 independent replications of sequences of future macro shocks.* Note that for comparability between the optimal policy *status quo* replacement policies, we used the same initial conditions in each simulation and in each simulation replication the 50 machines simulated under the *status quo* replacement policy were subjected to the same sequence of macro shocks as the corresponding fleet of 50 machines that were simulated under the optimal policy. Thus, our simulations provide a type of *controlled experiment* that would be very difficult to carry out in the real world.

The ratio of the present discounted profits to the OEC of each make/region combination can be considered as *rates of return* on the initial investment of an excavator at a specific location. For the example above of a make 4 excavator, the value to OEC ratio is 3.07, or expressed as a percentage rate of return this amounts to 207% — i.e. the initial investment of OEC required to buy the first excavator results in a total payoff in terms of expected future profits (on the current excavator and all future ones the fleet will replace it with) that is 3.07 times the initial OEC.

We generally find very high rates of return for all makes of excavators, and for all regions of the country. Thus, though excavators are expensive (i.e. their OEC can be $140,000 or higher) they also generate high gross rental profits, and as a result they are an attractive investment, earning rates of return on *initial* investment in the OEC of the first excavator of 200% or higher. In other words, the ratio of the value of expected discounted profits to the initial OEC for excavators is generally 3 or higher, and in some cases over 4 to 1.

The regions with the highest returns are B and D: this may partly reflect "locational rents" — i.e. higher rental rates, revenues, and higher time utilization of their excavators if their locations are in an area where there is a high demand for excavators and their locations are closer or more convenient for more customers. We also found that make 3 excavators earn higher returns than the other makes: i.e. they have higher value to OEC ratios than other makes of excavators. On the other hand, make 4 excavators generally produced the lowest rates of return.

Table 6 summarizes the average *undiscounted profits* earned per machine in the various simulations we did for different regions and makes of excavators. Comparing to the infinite horizon simulations, we generally find *larger percentage increases in profits* in our simulations than from our infinite horizon calculations of the value functions in table 5. For example for the make 4 excavator in region A, table 5 predicts a 3.5% gain in expected discounted profits over an infinite horizon, whereas table 6 predicts that the average per machine *undiscounted profits* increase by 11.7% over our 5 year simulation horizon.

Precisely how does the optimal replacement policy result in profit increases? Table 4 already provided some insights into this: the optimal policy involves keeping excavators longer than the firms keep them under their *status quo* replacement policy, and though the older machines earn less revenue and have higher maintenance costs, the reduction in replacement costs more than makes up for these other factors, resulting in higher profits overall. This same pattern is repeated for other makes of excavators and for other regions. We generally see that the optimal replacement policy involves keeping machines *longer* than under the *status quo*. For example for the make 4 excavator in region A, the simulated mean age of replacement under the optimal replacement policy is 76.8 months which is nearly 50% higher than the mean age at which these excavators are replaced in region A under its *status quo* replacement policy. The standard deviation replacement ages is much larger under the *status quo* than it is under the optimal replacement policy (34.1 versus 8.4).

We find that the main gain in profits comes from a significant reduction in *replacement costs* which is

Table 6: Gain in undiscounted profits: excavators simulated for 60 months (20 replications)

Make/Region	A	B	C	D	E	F
Make 1	20.6%	4.8%	3.3%	12.5%	6.4%	6.8%
Make 2	15.7%	1.4%	2.9%	6.3%	0.9%	2.9%
Make 3	9.4%	1.5%	4.0%	4.0%	1.3%	1.4%
Make 4	11.8%	1.5%	2.4%	3.5%	1.4%	-0.2%
Make 5	23.9%	2.8%	3.6%	12.4%	2.9%	11.5%

achieved in part by increasing the age when the excavators are replaced, but there are additional gains due to *strategic timing of replacement over the business cycle.* Because expected revenues decline with age and maintenance costs increase with age, keeping machines longer before replacing them decreases profits somewhat, but this decline in profits is more than compensated by a much larger decrease in replacement costs. For example, for the make 4 excavators in region A, our simulations reveal that average undiscounted revenue per excavator over the 20 replications over a 60 month simulation horizon decrease by 3% to $293,320 under the optimal policy from $301,712 under the status quo. Maintenance costs increase by 57%, from $11,264 to $17,749 per machine. However average replacement costs fall by nearly 50%, from $80,837 under the *status quo* to only $41,317 under the optimal replacement policy. The fall in replacement costs overwhelms the fall in current operating profits (i.e. monthly revenues less maintenance costs) due to the fact that average age of the fleet increases under the optimal replacement policy, so overall profits (net of replacement costs) increase by nearly 12%.

Figure 6 shows the simulated paths of replacements and average fleet age for a single replication of the simulations we did for the make 4 excavator owned in region A. The dashed black lines in each panel illustrate the simulated sequence of macro shocks in this particular simulation replication. We see that replacements tend to be more *clustered* under the optimal replacement policy than under the *status quo* and this is largely because of the idiosyncratic factors driving replacements under the *status quo.* Note that replacements occur in this simulation only in the boom states (i.e. when the dash line is at its highest point indicating boom macro states) whereas replacements occur in smaller, more regular bunches under the *status quo* replacement policy. This reflects the *strategic timing* motive, which in the case of excavators dictates replacing older machines in boom periods to take advantage of the lower replacement costs in these states.

Why are replacement costs lower in boom states? Recall the top left panel of figure 4 which shows

Figure 6: Simulation of fleet of make 4 excavators in region A

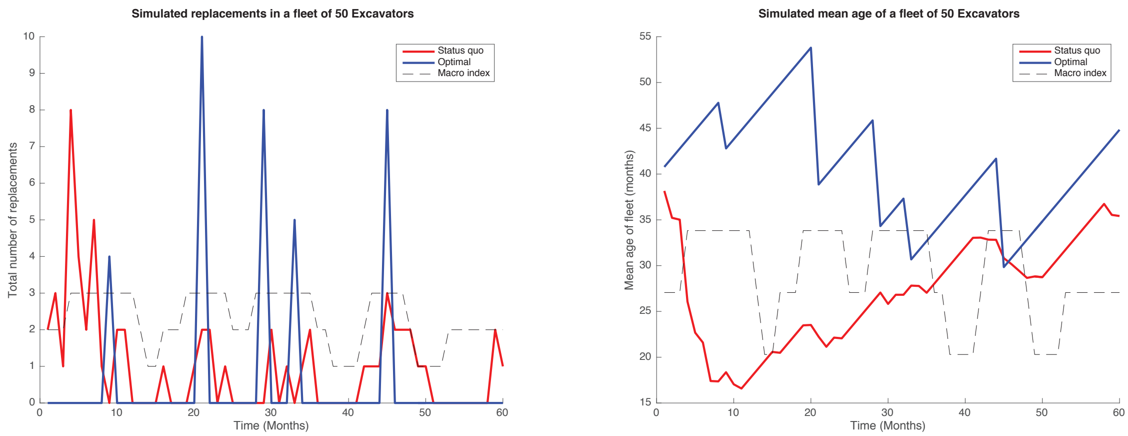

the predicted OEC and resale prices from our econometric model. Our econometric model predicts that OEC prices are lower in boom months, whereas resale prices of excavators that are 40 months or older are higher in boom months. Together these two effects imply that the cost of replacing an older make 4 excavator is lower in a boom month compared to a normal or bust month. So in effect, our dynamic programming algorithm has "discovered" a profitable way of *strategically timing replacements of machines to take advantage of predictable movements in machine prices over the business cycle.* In this case, the dynamic programming algorithm determines that the optimal replacement cycle is in fact *procyclical.*

Note that the concentrated nature of replacements under the optimal replacement policy is partly for strategic reasons but also due to the fact that we are simulating a model that does not allow for idiosyncratic shocks. Replacements will be less concentrated in specific months in a simulation of a model that allows for idiosyncratic shocks. However we did not make any assumptions about *quantity discounts* that may be available if a fleet manager undertakes *block replacements* of machines. If such discounts are available, it is possible to modify our model to allow the fleet manager to take advantage of them, and this can result in greater clustering of replacements.

6.2 Telescopic Booms

We complete our discussion of results by describing our findings for telescopic booms. Table 7 presents the gains in expected discounted values of profits for the three makes of booms in our data in one of the

Table 7: Gain in discounted profits from optimal replacement policy, telescopic booms

Make/Region	A	B	C	D	E	F
Make 1	14.2%	14.7%	13.1%	14.8%	25.9%	10.7%
Make 2	10.4%	10.0%	8.9%	10.6%	18.9%	8.2%
Make 3	10.3%	9.9%	8.9%	10.4%	17.6%	8.4%

company's locations in region A. Telescopic booms generate less profit per machine than excavators, but they are also about one third as expensive. Booms generate more profit per machine than skid steers, ranging from \$70,000 to about \$112,000 in discounted profit depending on the make and fleet, though we can see from Table 1 that average OEC for telescopic booms is more than twice as high as skid steers. Our calculations reveal that they expected discounted profits from skid steers range from \$50,000 to about \$95,000, depending on the make and region. Thus, we calculate that the overall rate of return from an initial investment in a telescopic boom to be lower than the corresponding rates of return for either excavators or skid steers.

On the other hand, we find that the percentage gains from adopting an optimal replacement policy are greater for telescopic booms than for excavators. The gains for booms range from an 8.2% gain in profits for make 2 machines in region F to a gain of 25.9% gain for make 1 machines in region E. However as we see in table 5, the percentage increase in discounted profits from adopting an optimal replacement policy for excavators range from a low of 1.2% for make 2 machines in region B or C to a high of 10.4% for a make 5 excavator in region A. However we found even higher percentage gains in discounted profits from adopting the optimal replacement policy in the case of skid steers: the gains ranged from a low of 13.4% for a make 2 skid steer in this location in region C to a high of 59.2% for a make 1 skid steer in a similar location in region E.

The company's locations in region C appear to generate the highest profits for all makes of telescopic booms under both the *status quo* and optimal replacement policies. We also see that make 3 machines generate the highest per machine profits across all six regions in our data set, with the exception of region E which our model predicts earns the highest profits from it make 1 telescopic booms.

Figure 7 illustrates the optimal replacement policy for make 1 telescopic booms at one of the company's locations in region A. Here we clearly see that the optimal replacement policy is *countercyclical.* For example the first panel shows the replacement thresholds by month for the three key cases, i.e. where the month is bust month (blue line), normal month (black line) or boom month (red line). Thus, if June is a bust month, the optimal replacement rule entails replacing a make 1 telescopic boom that is older than about 38 months. However if June is a normal month, then the fleet manager should not replace any booms unless they are over 140 months old, and if the economy is in a boom, it should not replace any machine unless it is over 158 months old.

When we compare this replacement policy to what the company does in region A (the dashed S-curves in the middle panel of figure 7) we see that company follows a *pro cyclical replacement strategy.* That is, it is significantly more likely to replace an older telescopic boom *in a boom month relative to a normal or bust month.* For example, the median age of replacement in a boom month under the company's *status quo* replacement policy is predicted to be about 110 months old, versus the 158 month threshold under the optimal policy. In a normal or bust month, the median age of replacement under the *status quo* is about 138 months, which is just slightly less than the 140 month replacement threshold under the optimal policy.

Thus, in a normal month the fleet manager replaces its make 1 machines at about the same median age as the replacement threshold calculated for the optimal replacement policy, but in a boom month, the optimal policy implies keeping machines much longer than what the fleet manager does, and in a bust month the optimal policy prescribes the fleet to replace its make 1 machines when they are relatively new (i.e. only 38 months old) whereas the under the *status quo* the fleet does not start to replace its make 1 machines until they are at least 60 months old and even for machines that are 140 months old, there is only a 60% probability they will be replaced by the fleet manager under its *status quo* replacement policy.

We can see the countercyclical nature of the optimal replacement policy clearly in figure 8, which shows a single simulation replication of a fleet of 50 make 1 telescopic booms in region A. The top panel shows that there are two major replacement spikes at 14 and 50 months into the simulation. These spikes were triggered by the fact that the economy went into a bust period, as indicated by the dips in the dashed line indicating that the macro shock reached its lowest possible value ($s = 1$, indicating a bust month) in the 14th and 50th months of the simulation. These were relatively short bust periods of only 3 to 4 months in each case, and we also see small subsequent replacement spikes at 16 months and 53 months, respectively where some additional machines were replaced under the optimal policy. Otherwise the economy had been only in normal and boom months, and so the optimal replacement policy for region A entailed no replacements in those other months of the simulation. However there were a small number of replacements driven by idiosyncratic shocks under the *status quo* replacement policy throughout the 60 month simulation

Figure 7: Optimal replacement policy for make 1 telescopic booms in region A

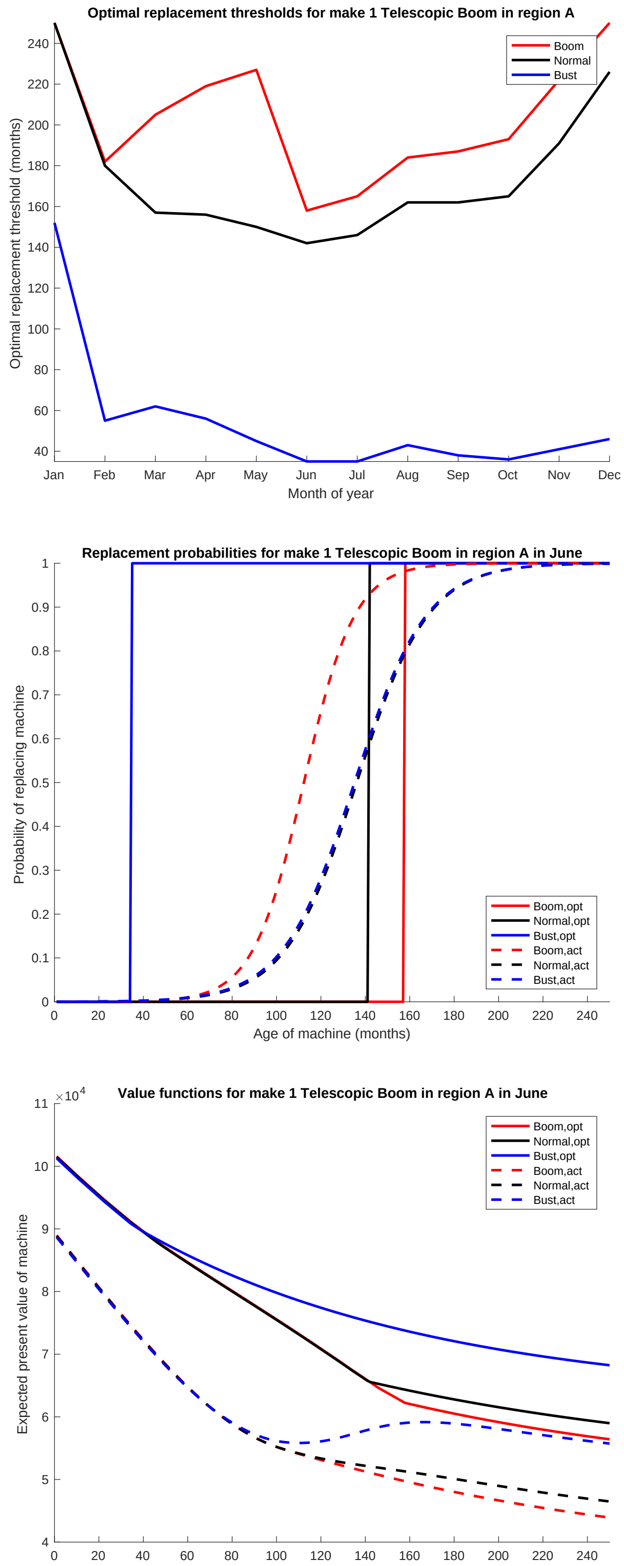

Figure 8: Simulation of fleet of make 1 telescopic booms in region A

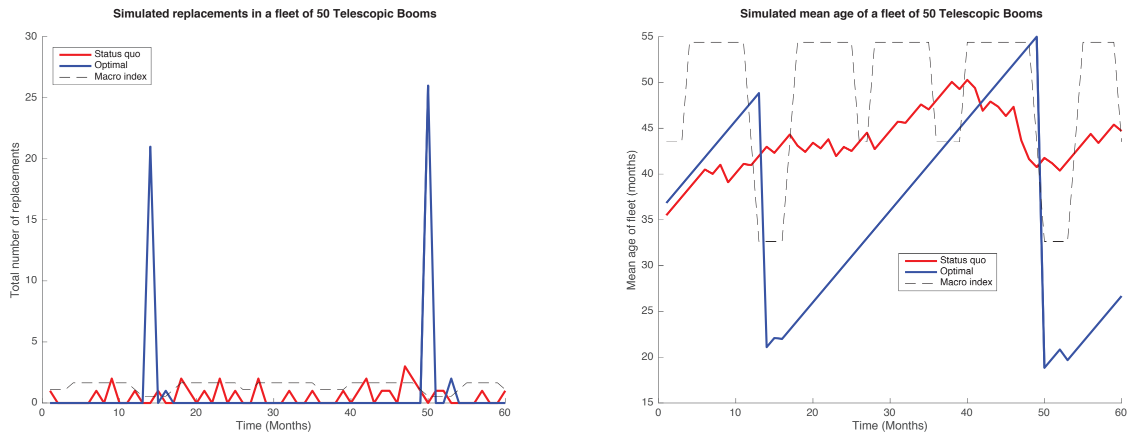

horizon.

The middle panel of figure 8 shows that the average age of the fleet increases linearly until month 14, whereas the average age of the fleet under the *status quo* increases less rapidly due to the rejuvenating effect of the replacements that occurred under the *status quo* policy. But by the time that the bust hits in month 14, the average age of the fleet under the optimal replacement policy is nearly 50 months old, and this exceeds the 38 month optimal replacement threshold in a bust month, so the firm ends up replacing 21 machines in its fleet that are over 38 months old in month 14 of the simulation.

The economy recovers from the short three month bust period and so there are no further replacements under the optimal policy until month 50 of the simulation when the macro state returns to the bust state. This motivates the firm to again replace all make 1 telescopic booms that are over 38 months old, and it replaces another 2 machines in month 53 of the simulation, bringing the average age of the fleet down to 20 months old, considerably younger than the average age of the fleet under the *status quo* which is about 45 months at that point.

We conclude our analysis of results with table 8, which presents the average undiscounted profits earned under the *status quo* versus the optimal policy over the 20 replications run over the 60 month simulation horizon. Here we find that the average gains from adopting the optimal replacement policy range from -2.1% (i.e. a loss) for make 2 booms in region F to a high of a 47% gain for make 1 booms in region E. Looking across makes of machines we find that the make 3 booms are the most profitable of the

Table 8: Gain in undiscounted profits: telescopic booms simulated for 60 months (20 replications)

Make/Region	A	B	C	D	E	F
Make 1	23.8%	23.2%	19.6%	26.3%	47.2%	6.5%
Make 2	13.8%	17.4%	8.5%	9.2%	27.4%	-2.1%
Make 3	6.9%	15.9%	6.4%	8.6%	30.4%	5.1%

three makes under the optimal policy, and looking across fleets we find higher profits for all three makes of telescopic booms in region C.

7 Conclusions

This paper has developed an *exploratory framework* for evaluating the profitability of rental machine replacement decisions that firms in our data set have actually made by comparing them with counterfactual "optimal" replacement decisions obtained from a *mathematical model* of optimal replacement decisions where the objective is to find a replacement strategy that *maximizes expected discounted gross profits (net cash flows) from the replacement of existing machines by new machines of the same make.*

We emphasize that there are many simplifying assumptions and abstractions from reality that we made to solve our mathematical model, and some of these assumptions may not be sufficiently good approximations to the very much more complex reality that actual firms face on a day to day basis.

Thus there is a distinct possibility that we are analyzing an *overly simplified model of the real problems that rental companies are facing* and for this reason the predictions of our model should be taken with a grain of salt until further explorations determine whether our model is a "sufficiently good approximation to reality" — or not.

We also acknowledged that the predictions of our model depend critically on the quality and accuracy of the *data* we were provided, and on the *econometric models and assumptions* we made to predict the rental revenues, maintenance costs, new OEC prices and used machine resale values that firms can expect to realize in their day to day operations. If the data are not the best, or our econometric model is providing misleading forecasts due to failure to use the best econometric techniques, then the conclusions we draw from our analysis are suspect from the standpoint of *providing reliable advice to rental companies on the replacement policies that will enable them to earn the highest possible profits.*

We have recognized that the overall problem that rental companies face is a very difficult, high dimensional problem that we do not have the data or the mental/computational resources to try to "crack" at the present time. However we followed a common sense approach known as *decomposition* in the computer science literature. That is, we decomposed the overall profit maximization problem that rental companies face and focused on the *replacement subproblem.* We argued that we can take other aspects of firm behavior/strategy as "given" (such as their policies over how they set rental rates and allocate lot space and capital to ownership of different types, makes, and models of machines) and focus on a simpler subproblem that we believe these companies need to solve, at least for their set of "core machines". We argued that firms can alter the solution to their replacement subproblems without significantly constraining or altering the other aspects of their fleet management and rental rate policies, especially because we have found that *rental rates are flat* and thus changing the mean age of replacement (at least when the changes are not sufficiently different from when firms currently replace machines) should not affect rental *rates,* though it does affect rental *revenues,* maintenance costs, and replacement costs.

We have shown that in some cases the optimal replacement policies calculated by our dynamic programming algorithm were not able to result in a significant improvement in profit relative to the firm's *status quo* replacement policies (such as for certain makes of excavators in some locations). In these cases the optimal replacement policies we calculated are similar to the replacement policies this firm is already using, so our analysis does not suggest any concern about the potential suboptimality of company decision making in these specific cases.

However in the majority of cases we studied, including most makes of excavators at most of the locations we analyzed, our calculations and simulations suggested that the firm could significantly improve its profitability by its their replacement policy, and a major share of the discounted increases in profits can be achieved over a relatively short horizon, i.e. within 5 years. Further, the optimal policies that we calculated numerically were significantly different from the policies this firm is using. For excavators, we found that in most of the locations and for most of the machine makes we analyzed, this firm could significantly increase their profitability by keeping its excavators *longer* than it does under the *status quo,* whereas for the other four types of machines we studied, we generally found significant profit gains from replacing the machines *significantly sooner* than they are replaced them under the firm's existing replacement policy. Indeed, for one of the regions and machine types — scissor lifts in region E — we found *huge gains* from replacing scissor lifts much sooner than this fleet is currently doing (such as after only 23 months for make 3 scissor lifts, compared to 114 months under the *status quo*). Moreover our calculations suggest that in location E the firm is keeping its scissor lifts past the *breakeven age* — that is, beyond the age where expected maintenance costs exceed expected rental revenue. This means that the firm is actually incurring an expected loss for every month beyond the breakeven age that it keeps its older scissor lifts in region E.

The other major difference between our calculated optimal replacement policies and the ones the fleet managers are using are that *the optimal replacement strategies are more sensitive to macro shocks.* For some machines such as excavators, we found that due to the way OEC and resale prices shift over the business cycle, replacement costs are lower in booms and thus the profit-maximizing replacement strategy involves doing most replacement investment in boom periods, so the optimal replacement policy is *pro-cyclical* for excavators. However for telescopic booms we found that OEC is significantly lower in bust periods, and thus it is significantly cheaper to replace telescopic booms in bust periods relative to normal or boom periods. Since replacement costs are such a significant part of the costs of the rental business, this implies that the *optimal replacement policy for telescopic booms is counter-cyclical.*

We will need to get more data and extend our model and subject it to further "stress tests" before we are willing to conclude that these preliminary findings are robust, generalizable findings. Further, we have ignored a very important consideration: *state-dependent borrowing constraints.* If many rental companies have lower ability to borrow or if the capital/retained earnings available to finance replacement investment "dries up" or becomes significantly more expensive during bust periods, then this is a consideration that we ignored in our analysis that could explain why the fleet managers in our sample follow a pro-cyclical rather than counter-cyclical replacement policy for telescopic booms and the other four machine types we analyzed. The firm may realize there is a profit opportunity from following a "contrarian" investment strategy, but once it takes borrowing constraints into account, it may conclude that it is infeasible to exploit this potential profit opportunity.

This suggest that we may need to further examine and try to relax some of the simplifying assumptions we made in this analysis, such as our assumption that the firm in our sample did not face borrowing constraints, particularly a drying up of capital for investment during bust periods. A related issue is the opportunity cost of capital. A countercyclical replacement investment strategy amplifies the procyclicality of the cash flow stream since it causes big cash flow outlays during bust periods, in return for cash inflows in normal and boom periods. Capital markets may strongly reward rental companies whose dividend payouts are more countercyclical, making their stocks better hedges for portfolio purposes, which could in turn lower their cost of capital. It may be important to try to account for the cyclical patterns of cash flows and the way stock markets (for publicly owned firms) or personal financial needs (for privately owned firms) lead firms to prefer cash flow streams that are less pro-cyclical and thus serve as a better hedge to overall comovements between the stock market and the business cycle.

If these "capital market" considerations are important, it suggests that different firms may have different optimal decision rules for similar assets. It may be that publicly traded firms where investors can hedge and diversify on their own and which have better access to capital markets can more easily pursue counter-cyclical strategies if the pro-cyclicality of the induced cash flow and dividend streams are not heavily penalized in the stock market. However capital constraints may be a real problem for smaller privately held firms. If so, the risk and "dividend smoothing" preferences of each firm might need to be accounted for in the model so that decision rules reflect those preferences. We believe it may be possible to take these capital market considerations into account and develop customized decision rules that reflect firms' practices, financial constraints, and preferences. If we can accurately capture these features in future versions of our models, the asset replacement rules that we derive from them will likely provide better guidance to firms.

A final area for future extensions is to allow for for a wider range of actions besides replacing an existing machine with a new one. This could include replacing existing machines with other used machines acquired at auction, or strategies that involve asset rejuvenation or rebuilds as a substitute to outright replacement. Rejuvenation or rebuilding adds to the useful life of the asset but at a cost. Therefore, there is a decision to be made about whether to rejuvenate an asset or simply dispose of the asset in favor of a new one or a different asset. There is also the question of how extensive the rejuvenation should be. The effect of rejuvenation on useful life and on reliability, or the probability of failure can also be included in the model and help produce more relevant, realistic and profitable decision rules.

References

Richard E. Bellman. *Dynamic Programming*. Princeton: Princeton University Press, 1957.

Sungjin Cho and John Rust. The flat rental puzzle. *Review of Economic Studies*, 77(2):534–559, 2010.

C. S. Chow and J. N. Tsitsiklis. The complexity of dynamic programming. *Journal of Complexity*, 5: 466–488, 1989.

John W. McClelland, Michael E. Wetzstein, and Richard K. Noles. Optimal replacement policies for rejuvenated assets. *American Journal of Agricultural Economics*, 71(1):147–157, 1989.

Harry J. Paarsch and John Rust. Valuing programs with deterministic and stochastic cycles. *Journal of Economic Dynamics and Control*, 33-3:614–623, 2009.

John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033, 1987.

John Rust. Using randomization to break the curse of dimensionality. *Econometrica*, 65-3:487–516, 1997.

Paul A. Samuelson. Some aspects of the pure theory of capital. *Quarterly Journal of Economics*, 5: 469–496, 1937.

Herbert A. Simon. *Administrative Behavior: a Study of Decision-Making Processes in Administrative Organization*. New York: Macmillan, 1947.